arXiv:1608.02926v2 [cs.CL] 26 Aug 2016

A pragmatic theory of generic language

Michael Henry Tessler (mtessler@stanford.edu)
Noah D. Goodman (ngoodman@stanford.edu)
Department of Psychology, Stanford University

Running head: PRAGMATIC THEORY OF GENERIC LANGUAGE

CORRESPONDING AUTHOR
MICHAEL HENRY TESSLER
STANFORD UNIVERSITY, DEPARTMENT OF PSYCHOLOGY
450 SERRA MALL, BLDG. 420, RM. 316
STANFORD, CA 94305
MTESSLER@STANFORD.EDU
757 - 561- 7071

Abstract

Generalizations about categories are central to human understanding, and generic language (e.g. *Dogs bark.*) provides a simple and ubiquitous way to communicate these generalizations. Yet the meaning of generic language is philosophically puzzling and has resisted precise formalization. We explore the idea that the core meaning of a generic sentence is simple but underspecified, and that general principles of pragmatic reasoning are responsible for establishing the precise meaning in context. Building on recent probabilistic models of language understanding, we provide a formal model for the evaluation and comprehension of generic sentences. This model explains the puzzling flexibility in usage of generics in terms of diverse prior beliefs about properties. We elicit these priors experimentally and show that the resulting model predictions explain almost all of the variance in human judgments for both common and novel generics. We probe the theory in more detail, and find that generic language depends in a fundamental way on subjective beliefs, not mere frequency. This theory provides the mathematical bridge between the words we use and the concepts they describe.

Disclosures and Acknowledgments

This research was supported in part by National Science Foundation Graduate Research Fellowship DGE-114747 (to M.H.T.), by John S. McDonnell Foundation Scholar Award 220020252 (to N.D.G.) and Office of Naval Research Grant N00014-13-1-0788 (to N.D.G.) Funding agencies had no other role other than financial support.

All authors contributed in a significant way to the manuscript. M. H. Tessler and N. D. Goodman developed theory and the study concept and design. M. H. Tessler performed research and analyzed data. M. H. Tessler and N. D. Goodman wrote the paper. All authors have read and approved the final manuscript.

Authors declare no conflict of interest.

Most would agree that *Swans are white*, but certainly not every swan is. This type of utterance conveys a generalization about a category (i.e. SWANS) and is known as a generic utterance (Carlson, 1977; Leslie, 2008). Communicating generically about categories is useful because categories themselves are unobservable (Markman, 1989). It is believed that every language can express generic meaning (Behrens, 2005; Carlson & Pelletier, 1995), and that generics are essential to the growth of conceptual knowledge (Gelman, 2004) and how kinds are represented in the mind (Leslie, 2008). Generic language is ubiquitous in everyday conversation as well as in child-directed speech (Gelman, Goetz, Sarnecka, & Flukes, 2008), and children as young as two or three understand that generics refer to categories and support generalization (Cimpian & Markman, 2008). Additionally, generics are the primary way by which speakers discuss social categories, making them key to propagating stereotypes (Gelman, Taylor, Nguyen, Leaper, & Bigler, 2004; Rhodes, Leslie, & Tworek, 2012; Leslie, Cimpian, Meyer, & Freeland, 2015) and impacting motivation (Cimpian, 2010). Despite their psychological centrality, apparent simplicity, and ubiquity, a formal account of generic meaning remains elusive.

The major issue in formalizing generic language is determining what makes a generic sentence true or false. At first glance, generics feel like universally-quantified statements as in *All swans are white*. Unlike universals, however, generics are resilient to counter-examples (e.g., *Swans are white* even though there exist black swans). Intuitions, then, fall back to something more vague like *Swans, in general, are white* because indeed most swans are white. But mosquitos, in general, do not carry malaria, yet everyone agrees *Mosquitos carry malaria*. Indeed, it appears that any *truth condition* stated in terms of how common the property is within the kind violates intuitions. Consider the birds: for a bird, being female practically implies you will lay eggs (the properties are present in the same proportion), yet we say things like *Birds lay eggs* and we do not say things like *Birds are female*.

How particular interpretations arise from generic language is also a mystery. The case of *Mosquitos carry malaria* suggests the generic must in some way be analogous to "some" (i.e. *Some mosquitos carry malaria.*). Yet, the empirical literature suggests generics are often interpreted as implying the property is widespread: There is a difference between *Some swans have hollow bones* and *Swans have hollow bones* (Gelman, Star, & Flukes, 2002). When interpretations are compared directly against truth conditions (i.e., how prevalent a property would need to be for the generic to be true), both children and adults infer the property is *more* prevalent when told the generic than when judging the generic true (Cimpian, Brandone, & Gelman, 2010; Brandone, Gelman, & Hedglen, 2014), suggesting that speakers who use generics can exaggerate evidence to their listeners.

How can generics have such flexible truth conditions while simultaneously carrying strong implications? In this paper, we explore the idea that the core meaning of a generic statement is simple,

but underspecified, and that general principles of communication may be used to resolve precise meaning in context. In particular, we develop a mathematical model that describes pragmatic reasoning about the degree of prevalence required to assert the generic. We find that this formalism resolves the philosophical and empirical puzzles.

In Section 1, we review extant theories of generics from the linguistic and psychological literatures. In Section 2, we describe our pragmatic theory of generic language in formal terms and illustrate the model through simulation. In Section 3, we conduct a set of experiments to test the *truth conditions of generic statements*, and find our model predicts the patterns in human endorsement of familiar generic sentences. We also find that the prior beliefs used in the Bayesian model reflect conceptual structure, and this provides an understanding of the conceptual implications of generics. In Section 4, we conduct a set of experiments to test *interpretations of novel generic sentences*, again finding our model predicts human judgments with high quantitative accuracy. In Section 5, we investigate the underlying semantic scale in more detail, and find that the core meaning of generic language depends in a crucial way on subjective beliefs, not mere frequency.

The philosophical puzzle of generic language: Truth conditions

Generics express a relation between a kind K (e.g., ROBINS) and a property F (e.g., LAYS EGGS), such that the property can also be said to be applicable of an individual (i.e., the bird in my backyard lays eggs). Bare plural statements (e.g., *Robins lay eggs*) tend strongly to yield a generic meaning (Carlson, 1977), though other forms can express such a meaning sometimes (e.g., *A mongoose eats snakes.*).

Given that generics express a property that can be applied to individuals, it would seem intuitive that the number of individuals with the property would be what makes the statement true or false. Counter-examples like *Mosquitos carry malaria* and *Birds lay eggs* v. *Birds are female* stifle such intuition. Semantic theories that appeal to the statistics of the world (i.e., how many Ks have F) try to rescue the notion that a generic expresses something like *Ks, in general, have F*, where *in general* is often either restricted to particular domains or is calculated with particular distortions. Alternative, *conceptually*-based theories take a generic to mean *Ks, in virtue of being the kind of things that Ks are, have F*. These accounts appeal to structured, conceptual representations in deciding what kinds of generic statements are true or false. Statistical and conceptual theories express the major contrasting views of the truth conditions of generic statements (Carlson, 1995).[1]

[1]We use the terms statistical and conceptual to refer to what Carlson (1995) referred to as "inductive" and "rules and regulations" views, respectively.

Statistical accounts of generics

Statistical accounts take the **property prevalence** to be fundamental: *Birds lay eggs* means *Birds, in general, lay eggs*. Of course, birds do not in general lay eggs (it's only the adult, female ones that do). The primary way of dealing with such issues is to posit domain restrictions ("implicitly, we are only talking about the females") when there are "salient partitions" (Carlson & Pelletier, 1995). The most fully-developed theory on this front is due to Cohen (1999). Let's first introduce some notation:

For a given kind K (e.g. ROBINS) and property F (e.g. LAYS EGGS), we refer to the probability that an instance of kind K has property F, that is $P(F \mid K)$, as the *prevalence* of F within K. Logical quantifiers can be described as conditions on prevalence (i.e. *some* is $P(F \mid K) > 0$, *all* is $P(F \mid K) = 1$). Assuming the generic relates to the property prevalence, then the simplest meaning would similarly be a threshold on prevalence: $P(F \mid K) > \theta$. Cohen (1999) takes $\theta = 0.5$; that is, *Robins lay eggs* is roughly taken to mean *More than half of (relevant) robins lay eggs*.

Cohen introduces constraints into the computation of prevalence: $P(F \mid K)$. In particular, prevalence is computed with respect to a *partition set*. For example, the property *lays eggs* induces a set of alternatives that all have to do with procreation (e.g. *gives birth to live young*, *undergoes mitosis*, ...). The individuals that enter into the prevalence computation are only those individuals that would satisfy one or another alternative. That is, the only individuals under consideration are the female members of kinds because only the female members can plausibly satisfy one or another of the alternatives. The inference that this sort of domain restriction is applied, as well as other constraints posited in the theory, are thought to be a pragmatic inference, though specifying the details of the pragmatics is beyond the scope of Cohen's theory. Without a well specified theory of pragmatics, however, we have no general solution to map from property prevalence to truth judgment. Conceptual information seems to be behind the pragmatic inference, but how Cohen's and other prevalence-based accounts relate to such knowledge remains obscure (Carlson, 1995). [2]

Conceptual accounts of generics

Conceptual accounts of generics emphasize the structure of generic knowledge (Prasada, 2000), and view generic utterances as the way of expressing special mental relationships between kinds and properties (Leslie, 2008; Prasada, Hennefield, & Otap, 2012). From this perspective, *Bishops move diagonally* because those are the rules of the game, not because they *tend to* move diagonally. How do we come to know the rules of the game, though?

Leslie (2007)'s influential theory posits that generics are tied to a "default mode of generaliza-

[2] However, see Cohen (2004) for a discussion of how his semantic constraints relate to different kinds of generics and different kinds of conceptual representational frameworks found in cognitive science.

tion". This "default mode" comes equipped with the ability to single-out *striking properties* (e.g. properties which are dangerous or appalling) as particularly useful aspects of the world to know about. Hence, *Mosquitos carry malaria* is true because carrying malaria is striking, and thus, a useful bit of information to convey. The default mode can also distinguish "negative" counter-instances of a property (e.g. a bird that doesn't lay eggs) from "positive" counter-instances (e.g. a hypothetical bird that bears live young). Generics are much less reasonable when positive counter-instances exist. Hence, *Birds are female* seems weird because "being male" is a positive counter-instance of "being female", but since there are no birds that bear live young, *Birds lay eggs* is fine.

Parallel work in the psychology of concepts supports this perspective. Prasada and Dillingham (2006) and later Prasada, Khemlani, Leslie, and Glucksberg (2013) distinguish between *principled*, *statistical*, and *causal* relations within concepts. Generics like *Birds lay eggs* (in which only a minority of K have F) exhibit characteristics of principled connections (operationalized by endorsement of the phrase *In general, Ks have F*), supporting Leslie (2007)'s typology. Striking generics (e.g. *Mosquitos carry malaria*) show characteristics of *causal connections* (operationalized using the phrase *There is something about Ks that cause them to F*). The fact that generics about different properties license different kinds of inferences is taken as evidence that the generics themselves represent different kinds of relations. Statistical information takes a backseat to the conceptual structure.

The pragmatic perspective

We find both the statistical and conceptual accounts compelling, but an important perspective is missing. We note that generic language is not unique in its flexibility. Understanding language in general depends upon assumptions interlocutors make about each other and what content is under discussion. The pragmatic lens reveals that utterances carry a mosaic of interpretations with a complex sensitivity to context (Clark, 1996; Grice, 1975; Levinson, 2000). Can the puzzles of generic language be understood as effects of pragmatic reasoning? If so, we may be able to get away with a relatively simple, statistical, semantic theory. Abstract mental representations, then, may be the conceptual backdrop against which generic language is interpreted.

We pursue such a line of inquiry, assuming the simplest truth-conditional meaning: a threshold on prevalence $P(F \mid K) > \theta$ (cf. Cohen, 1999). No fixed value of the threshold, θ, would allow for the extreme flexibility generics exhibit (e.g. *Mosquitos carry malaria*; *Robins lay eggs* v. *Robins are female*), so instead we allow this threshold to be established in context through pragmatic reasoning. Such an inference would depend on background knowledge about properties and categories—potentially structured, conceptual knowledge. This inference, nonetheless, is a general

mechanism of understanding language, not specific to interpreting generic statements. We formalize this hypothesis in the Rational Speech Act (RSA) theory—a formal, probabilistic theory of language understanding. RSA is derived from social reasoning, formalized as recursive Bayesian inference between speaker and listener (Frank & Goodman, 2012; Goodman & Stuhlmüller, 2013; Goodman & Frank, To appear; see also, Franke, 2009; Franke & Jäger, 2015). Goodman and Frank (To appear) provides a good introduction to the RSA framework and Appendix A presents a brief tutorial on RSA for the reader unfamiliar.

We follow the treatment of RSA applied to vague adjectives (e.g. *tall*), where a semantic variable is left underspecified (Lassiter & Goodman, 2013, 2015). We model the interpretation of generic statements by a pragmatic listener (L_1), concerned with learning the prevalence of a certain property in a certain category, $x = P(F \mid K)$. She resolves the likely prevalence $x \in [0,1]$ upon hearing a generic utterance u, by integrating her prior beliefs about x with the generative process of the utterance – an informative speaker (S_1).

$$L_1(x, \theta \mid u) \quad \propto \quad S_1(u \mid x, \theta) \cdot P(x) \cdot P(\theta) \tag{1}$$

In addition to not knowing x *a priori*, the listener does not know the value of the prevalence threshold, θ, used to establish the meaning of the generic (as described below). Because the speaker's choice of utterance can depend on the threshold, the listener must reason about the threshold as well, based on her prior distribution over possible thresholds $P(\theta)$. In principle, thresholds could be learned over time for different contexts, but here we assume the listener has no informative knowledge about the semantic variable: $\theta \sim \text{Uniform}([0, 1])$.

To resolve an appropriate meaning, the listener uses her background knowledge about the property $P(x)$ and her intuitive theory of a speaker S_1 as a rational actor who has the goal of providing information to a hypothetical literal listener (L_0):

$$S_1(u \mid x, \theta) \quad \propto \quad \exp\left(\lambda_1 \cdot \ln(L_0(x \mid u, \theta))\right) \tag{2}$$

The listener believes the speaker S_1 knows the threshold θ. That is, the speaker had a particular meaning in mind, and the listener can use pragmatic reasoning to help resolve that meaning (Lassiter & Goodman, 2013, 2015; Goodman & Lassiter, 2015).

The proportionality in Eq. 2 implies normalization over a set of alternative utterances. We consider the simplest set of alternatives, including only the production of a generic utterance and the possibility of staying quiet (and we assume no difference in production cost between these alternatives). We adopt the simplest meaning of a generic utterance in terms of prevalence, a threshold function: $[\![\text{K F}]\!](x, \theta) = x > \theta$. The silent or *null* utterance alternative carries no information

(i.e. $[\![null]\!](x,\theta) = T$) and produces a posterior distribution in the listener identical with the prior.[3] The *null* utterance alternative gives the generic its communicative force; the option of staying silent makes the generic utterance into a *speech-act*. S_1 decides what is informative with respect to an idealized literal listener (L_0). The literal listener has access to the threshold θ as well, and simply restricts her prior beliefs to situations where the truth-functional denotation of the utterance, $[\![u]\!]$, is true (for the silent utterance: all worlds; for the generic utterance: the worlds with prevalence greater than θ).

$$L_0(x \mid u, \theta) \quad \propto \quad \delta_{\,[\![u]\!](x,\theta)} \cdot P(x) \tag{3}$$

The pragmatic listener L_1 (Eq. 1) is a model of generic interpretation: Upon hearing a generic, what prevalence is a listener likely to infer? This model assumes a threshold semantics, but doesn't require us to specify the threshold *a priori*: It is inferred with pragmatic reasoning. We can now imagine a speaker who reasons about this kind of listener:

$$S_2(u \mid x) \propto \exp\left(\lambda_2 \cdot \ln \int_\theta L_1(x, \theta \mid u)\, d\theta\right) \tag{4}$$

This speaker considers the thought-processes of a listener who understands generics (L_1, Eq. 1) and decides if the generic is a good (albeit, vague) way to describe the prevalence x. S_2's decision (like the simpler S_1's decision) is with respect to the alternative of saying nothing: He will choose to produce the generic when the true prevalence x is more likely under L_1's posterior than under her prior. Critically, in contrast to the hypothetical speaker S_1, speaker S_2 doesn't actually know what the generic means (i.e. doesn't have access to the threshold θ), but knows that it will be inferred by L_1, and integrates over the likely values she'll consider. In line with the Rational Speech Act theory, we assume this speaker is a soft-max optimal speaker, with degree of optimality governed by parameter λ_2. S_2 is thus a mapping from prevalence x to the probability of producing a generic utterance. Since the only alternative to producing the generic is to stay silent, S_2 is interpreted as a model of endorsement, felicity, or truth judgments (Degen & Goodman, 2014). We implemented this model in the probabilisitic programming language WebPPL (Goodman & Stuhlmüller, 2014), and a fully specified version of the model can be found at `http://forestdb.org/models/generics.html`

This model makes the prediction that the interpretation of a generic statement (and hence, it's corresponding truth judgment) depends on the prior on prevalence $P(x)$. We can explore the model by running simulations of L_1 and S_2 under various schematic priors. $P(x)$ is a distribution on the

[3] This alternative can be realized in at least two other ways: the speaker could have said the negation of the utterance (i.e. *It is not the case that Ks have F.*) or the negative generic (i.e. *Ks do not have F.*). All results reported are similar for these two alternatives, and we use the alternative of the *null* utterance for simplicity.

prevalence of a given property (e.g. LAYS EGGS) across animal categories. In Figure 1, are L_1 (Eq. 1) posterior prevalence distributions on x (red solid line) for several example prior prevalence distributions $P(x)$ (blue dashed line), as well as the S_2 generic endorsement probability for different levels of prevalence. We name these example prior distributions to suggest properties that might be associated with such priors. (Later, we will empirically measure these priors for properties of interest.) First, we explore generic interpretation and endorsement for priors of three different shapes (left column). For each of these, the L_1 posterior distribution (red solid) over prevalence is heavily driven by the prior (blue dashed). S_2 endorsement probabilities for the generic (black solid) increase as a function of prevalence, and what counts as "true" (in terms of the prevalence) depends on the prior.

To take one example, consider the distribution over what might be the property "are female" (top-left facet). The *a priori* prevalence is centered around 0.5. Because of pragmatics, the pressure to be *truthful* will drive likely threshold values below 0.5 (lower values are more likely to be true). At the same time, the pressure to be *informative* will drive the threshold values up: Higher values are more likely to be informative. The result is a posterior over prevalence that is only marginally greater than the prior, making higher prevalence values more likely after hearing the generic. But the relative information gain is very little (the posterior is not very different from the prior), and thus the S_2 model is reluctant to endorse the generic unless the property is exceedingly prevalent. The same basic phenomenon can be observed for the other two example priors (left column): The posterior over prevalence heavily depends upon the prior, but is also not very different from the prior.

We now imagine what would happen if there were some kinds where the property was completely absent, while being present at some rate in other kinds. Figure 1 (right column) shows this possibility: mixing the distribution to the left of it with a second component at 0% prevalence. Consider the schematic prior over "lay eggs". We see the pragmatic listener L_1's posterior prevalence distribution is not very different from the interpretation that doesn't include the second component in the prior (here, "are female"), but it is dramatically different from the prior with two components (compare red with blue dashed line for the right column). This suggests that when many categories have 0% prevalence, lower thresholds are informative. Indeed, the S_2 model predicts that the generic becomes felicitous at lower prevalence levels (compare black line of left v. right column). For "lay eggs", when the property is prevalent in 50% of the kind (e.g., 50% of birds lay eggs), endorsement of the generic (e.g., *Birds lay eggs*) by the S_2 is roughly 0.85; for "are female", with the same prevalence (50% of birds are female), endorsement for *Birds are female* is only 0.50: It is judged to be neither true nor false. The important result is the asymmetry: the first generic can be endorsed more strongly than the second, at the same prevalence level; the exact endorsement rates depend on quantitative aspects of the priors, which must be determined

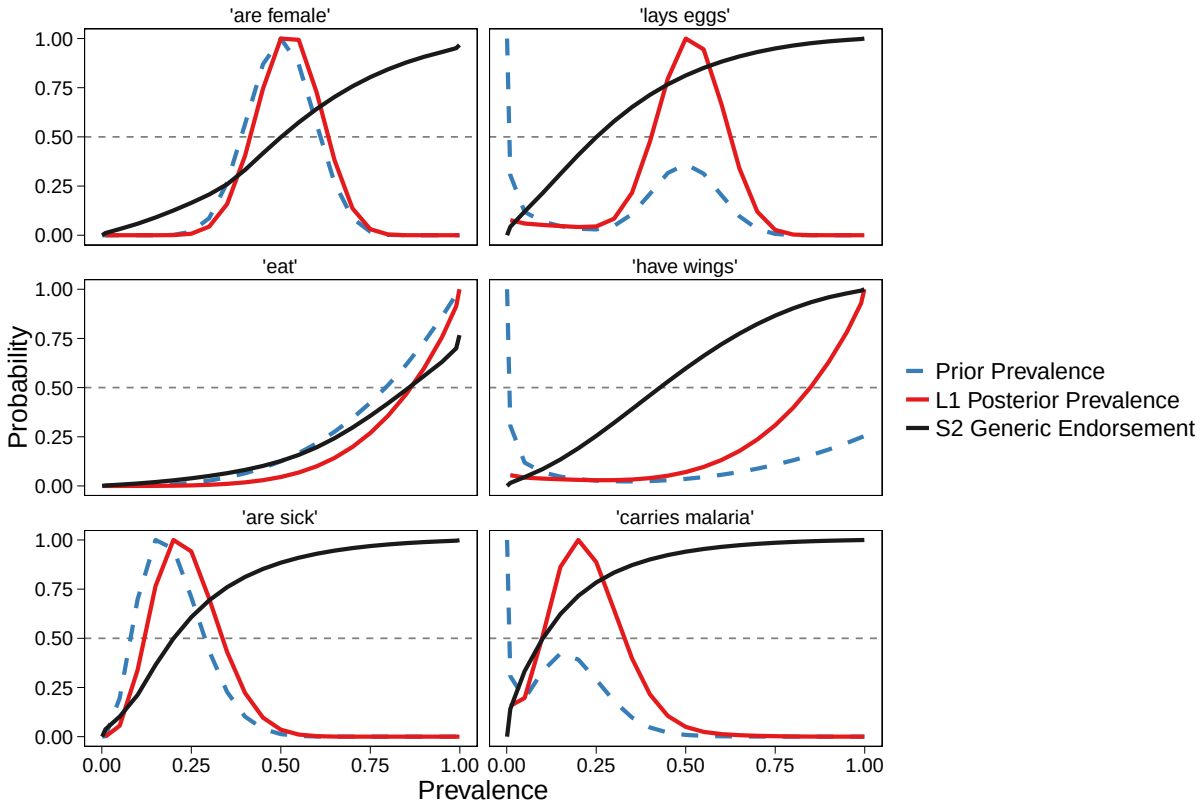

Figure 1: Schematic prior distributions for prevalence x (blue dashed), the pragmatic listener L_1 model's posterior distribution over prevalence upon hearing a generic utterance (red), and speaker S_2 model's endorsement of a generic utterance for different levels of prevalence (black). The names given to these priors are meant to be suggestive of what kinds of properties these distributions might correspond to. The left column uses prevalence priors modeled as Beta(15,15), Beta(4,1), and Beta(4,16) distributions. The right column uses a prior distribution that is a mixture of the distribution to the left of it with a second component, modeled as Beta(0.5, 4.5), reflecting categories with 0% property prevalence. Horizontal dashed line at 0.5 is for convenience of comparing the point at which an utterance becomes judged as more true than false for S_2. Note that the prior distribution over prevalence will be the same as L_1's posterior distribution upon hearing the "null" utterance.

empirically.

The model thus predicts differences in truth judgments depending on the prevalence prior. The first test of our theory, then, will be to see if these predictions correspond to human truth judgments of familiar generic sentences.

Empirical test 1: Flexible truth conditions

Any theory of generic language must explain their puzzling flexibility of usage with respect to prevalence. That is, *Mosquitos carry malaria* and *Birds lay eggs* are reasonable things to say, but *Birds are female* is not. The pragmatic speaker model S_2, Eq. 4, is a model of truth judgments. We test our model on thirty generic sentences that cover a range of conceptual distinctions discussed in the literature (Prasada et al., 2013): characteristic (e.g. *Ducks have wings.*), minority (e.g. *Robins lay eggs.*), striking (e.g. *Mosquitos carry malaria.*), false generalization (e.g. *Robins are female.*), and false (e.g. *Lions lay eggs.*). In additional to the canonical cases from the linguistics literature, we selected sentences to elicit the full range of acceptability judgments (intuitively, "acceptable", "unacceptable", and "uncertain") with low-, medium-, and high-prevalence properties.

The pragmatic speaker model S_2 is fully-specified except for the prior distribution over prevalence $P(x)$, which plausibly varies by the type of property in question. To compare the model to empirical truth judgments, we thus first measure the prior distribution over prevalence of these properties (Expt. 1a). In Expt. 1b, we collect human judgements about the acceptability of the generic sentences.

Experiment 1a: Prevalence priors

The prior $P(x)$ (in Eqs. 1, 3) describes the belief distribution on the prevalence of a given property (e.g. LAYS EGGS) across relevant categories. In exploring the model, we saw that the shape of this distribution affects model predictions, and this shape may vary significantly among different properties. We thus measured this distribution empirically for the set of properties (e.g. LAYS EGGS, CARRIES MALARIA; 21 in total) used in our target sentences.

Method

Participants

We recruited 60 participants over Amazon's crowd-sourcing platform Mechanical Turk (MTurk). Participants were restricted to those with US IP addresses and with at least a 95% MTurk work

approval rating (the same criteria apply to all experiments reported). 3 participants where unintentionally allowed to do the experiment for a second time; we excluded their second responses (resulting in $n = 57$). 2 participants self-reported a native language other than English; removing their data ($n = 55$) has no effect on the results reported. The experiment took about 10 minutes and participants were compensated $1.00.

Procedure and materials

On each trial of the experiment, participants filled out a table where each row was an animal category and each column was a property. In order to alleviate the dependence of the distribution on our animal categories of interest, participants generated half of the animal categories before viewing the properties; the other half were randomly sampled from a set corresponding to the generic sentences used in Expt. 1b (e.g. ROBINS, MOSQUITOS).

Participants began the experiment by seeing a list of 6 animal kinds and were asked to list 5 of their own. A column then appeared to the right of the animal names with a property in the header (e.g. "lays eggs"). Participants were asked to fill in each row with the percentage of members of each of the species that had the property (e.g. "50%"). Eight property–columns appeared in the table, and this whole procedure was repeated 2 times. In total, each participant generated 10 animal names and reported on the prevalence of sixteen properties for 22 animals (their own 10 and the experimentally-supplied 12). Properties were randomly sampled from a set of 21 properties associated with generics of theoretical interest, as described above. For a full list of the properties, and generic sentences used in Expt. 1b, see Table 2 (Appendix). The experiment can be viewed at `http://stanford.edu/~mtessler/experiments/generics/experiments/real-kinds/prior-2.html`.

Data analysis and model predictions

To process the priors data, we discretize the prevalence judgments to 12 discrete bins: $\{[0-0.01), (0.01-0.05), (0.05-0.15), (0.15-0.25), ..., (0.75-0.85), (0.85-0.95), (0.95-1]\}$, and look at the counts within each bin, after doing add-1 Laplace smoothing, as the relative probability of that prevalence. Using these priors, we can explore how $L_1(x, \theta \mid u)$, the pragmatic listener model, interprets a generic utterance (Figure 2a, insets). The prior beliefs over the prevalence of the property, $P(x)$, can also be interpreted as the pragmatic listener's posterior upon hearing the null utterance, because the null utterance has no information content. We see the interpretation of the generic is quite variable across our empirically measured priors. For instance, in the case of CARRIES MALARIA, the prior is very left-skewed; here, the threshold θ can plausibly be quite low while still being informative, since a low threshold still rules out many possible alternative kinds

(and their corresponding degree of prevalence). Properties like DOESN'T ATTACK SWIMMERS are very right-skewed; here, even a relatively high threshold would not result in an informative utterance (intuitively, as not many kinds would be ruled out), and so the generic is unlikely to be used by speaker S_2 unless the property is practically-universal within the target category. Some properties have priors that are unimodal with low variance (e.g. IS FEMALE); these properties are present in every kind in almost exactly the same proportion and thus are too obvious and certain to allow for an informative generic utterance: The posterior is not very different from the prior. With $P(x)$ now empirically established, we can test if our speaker model predicts human truth judgments of generic statements about these properties.

Experiment 1b: Truth judgments

Method

Participants

We recruited 100 participants over MTurk. 4 participants were excluded for failing to pass a catch trial. 5 participants self-reported a native language other than English; removing their data has no effect on the results reported. The experiment took about 3 minutes and participants were compensated $0.35.

Procedure and materials

Participants were shown thirty generic sentences in bare plural form, one after another. They were asked to press one of two buttons to signify whether they agreed or disagreed with the sentence (see Table 2 in Appendix for complete list). The thirty sentences were presented in a random order between participants and covered a range of conceptual categories described above. Approximately 10 true, 10 false, and 10 uncertain truth-value generics were selected.

As an attention check, participants were asked at the end of the trials which button corresponded to "Agree". 4 participants were excluded for failing this trial.

Data analysis and results

As a manipulation check, the first author assigned an *a priori* truth-judgment (true/false/indeterminate) to each stimulus item. This was a significant predictor of the empirical truth judgments: true generics were significantly more likely to be agreed with than the indeterminate generics

($\beta = 3.14; SE = 0.15; z = 21.5$), as revealed by a mixed-effect logistic regression with random by-participant effects of intercept. Indeterminate generics were agreed with *less* likely than chance ($\beta = -0.49; SE = 0.09; z = -5.3$) but significantly more than false generics ($\beta = 2.09; SE = 0.14; z = 14.5$).

From the prevalence prior data (Expt. 1a), we estimate participants' beliefs about the prevalence of a property *for a given kind* (e.g. the percentage of ROBINS that LAY EGGS; see green intervals on Figure 2a insets, and Table 2 in Appendix). As a simple baseline hypothesis, we first explore whether these prevalence values themselves predict generic endorsement (e.g. does the fraction of ROBINS that LAY EGGS predict the felicity of *Robins lay eggs*?). We find a little over half of the variance in truth judgments data is explained this way ($r^2 = 0.599$; MSE=0.065; Figure 2b, right). This is not surprising given that our stimulus set included generics that are true with high-prevalence properties (e.g. *Leopards have spots.*) and generics that are false with low prevalence properties (e.g. *Leopards have wings.*). However, large deviations from an account based purely on target-category prevalence remain: Generics in which the target-category has intermediate prevalence (prevalence quartiles 2 and 3: $20\% < prevalence < 64\%$), are not at all explained by prevalence within those categories ($r^2_{Q2,3} = 0.029$; MSE = 0.110).

The pragmatic speaker model, S_2 in Eq. 4, predicts an endorsement probability for a generic sentence, given prior beliefs about the property, $P(x)$, and a prevalence-level, x, within the kind-of-interest. (That is, S_2 provides a model of whether someone who knows x would say the generic to someone who doesn't, but shares prior $P(x)$; as discussed above we adopt this as a model of agreement judgments.) We use the empirically estimated within-kind prevalence as the x that the speaker S_2 is trying to communicate, and use the empirically measured priors from Expt. 1a as the listener's prior $P(x)$. The S_2 model is then fully specified after setting the speaker optimality parameters λ_1 and λ_2 (in Eqs. 2, 4). Rather than fitting the parameters, we inferred them and integrated over their plausible values using Bayesian data analysis (Lee & Wagenmakers, 2014). We put uninformative priors over these parameters, with a range consistent with previous literature using the same model class: $\lambda_1 \sim \text{Uniform}(0,20)$, $\lambda_2 \sim \text{Uniform}(0,5)$. We learn about the *a posteriori* credible values of our model parameters by collecting samples from MCMC chains of 10,000 iterations removing the first 5,000 iterations, using the Metropolis-Hastings algorithm.

The above analysis provides a single estimate for model predictions, but it is based on noisy empirical measurements of $P(x)$. In order to estimate the impact of this empirical noise on our model predictions, we resampled the prior data (with replacement) for 57 participants worth of data, discretizing and binning as we did above and then inferring parameters. This procedure (resample prior, discretize and bin, infer parameters) was repeated 500 times to bootstrap the model predictions. The Maximum A-Posteriori (MAP) estimate and 95% Highest Probability Density

(HPD) interval for λ_1 is 0.5 [0.004, 12.7] and λ_2 is 1.7 [1.3, 2.1]. [4]

We compare the model's posterior predictive distribution of generic endorsement to the empirical truth judgments. (The posterior predictive distribution marginalizes over the inferred parameter values to produce predictions about what the data *should look like* given the pragmatics model and the observed data. This is akin to fitting the parameters and is the critical step in model validation: It shows what data is actually predicted by the model.) As we see in Figure 2b, the pragmatic speaker model S_2, using empirically measured priors, explains nearly all of the variance in human truth judgments ($r^2 = 0.98$; MSE=0.003; Figure 2b).

Generics that received definitive agreement or disagreement are predicted to be judged as such by the model (corners of Figure 2b), including items for which target-category prevalence is not a good indicator of the acceptability (e.g. *Mosquitos carry malaria*, for prevalence quartiles 2 and 3, $r^2_{Q2,3} = 0.955$; MSE=0.005; Figure 2b, intermediate shades). We also see the generics truth judgment model predicts uncertain truth judgments: for instance, *Robins are female* is judged by both the model and human participants to be neither true nor false. *Sharks don't attack swimmers*, while true of most sharks, is judged to be not a good thing to say by both participants and the model. This is strong evidence that the puzzling flexibility of generic truth-conditions can be understood with a simple semantic theory coupled with basic communicative principles (*be truthful*, *be informative*) operating over diverse prior beliefs about the properties, all of which are at play in understanding language.

Extended analysis of priors: Conceptual structure

The pragmatics model only has a prevalence-based semantics yet it is able to explain the flexibility in truth judgments for a diverse range of generic statements. Conceptual accounts of generic statements have looked beyond prevalence, to structured knowledge representations as the critical factor in generic meaning (Leslie, 2007; Prasada et al., 2013). We can interrogate our formal model to see what is driving its predictions, and, in particular, we ask whether structured representations might effect our model after all.

For a given property, the prior distribution on prevalence $P(x)$ is a single distribution. However, the distribution may be structured as the result of deeper conceptual knowledge. For instance, if participants believe that some kinds have a causal mechanism that *could* give rise to the property, while kinds others do not, then we would expect $P(x)$ to be structured as a mixture distribution (cf., Griffiths & Tenenbaum, 2005). We know that prior knowledge plays a fundamental role in leading the pragmatics model to endorse or reject generics. We now explore the hypothesis that this is

[4] The fact that λ_1 is credibly less than 1 suggests that the generic utterance may be more costly than "staying silent" (our model assumes equal cost). We maintain using only the two λ parameters in the model for simplicity.

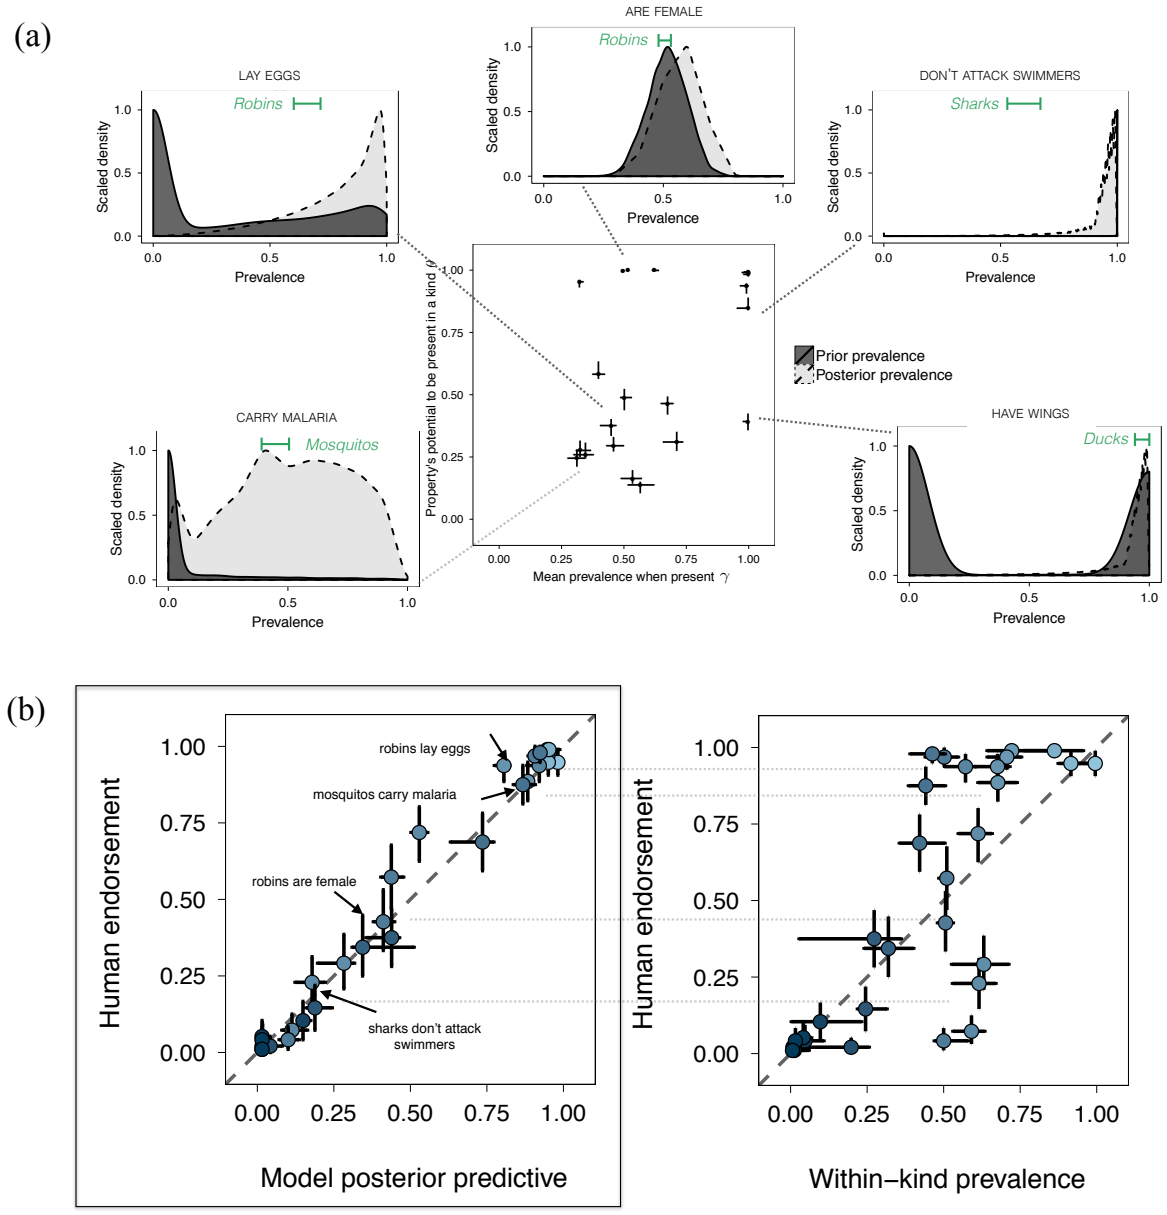

Figure 2: Endorsing familiar generics. (a) Prevalence prior distributions empirically elicited for twenty-one animal properties. Prior distributions summarized by two parameters of a structured Bayesian model: ϕ—-a property's potential to be present in a category—-and γ—-the mean prevalence when it is possible for the property to be present in a category. Inset plots display example empirical prior distributions over prevalence and corresponding L_1 model predictions: the posterior after hearing a generic utterance. Intervals on the top of insets show human judgments about the prevalence of the property within a target category. (b) Human acceptability judgments compared with model predictions (left) and the target-category prevalence (right) for thirty generic utterances about familiar animals and properties. Color denotes target-category prevalence of the property, with lighter colors indicating higher prevalence. Error bars denote 95% Bayesian credible intervals.

at least partly because these priors are structured into two components: kinds that *can* have the property and other kinds that *cannot*. We explore this possibility by formulating a mixture model for the prevalence priors, and exploring how well it fits the prior data elicited in Expt. 1a.

Data analysis

If a kind can have the property, we assume the prevalence follows a Beta distribution with mean γ and concentration ξ. If a kind cannot, we assume the prevalence is a Delta distribution, with all probability mass at 0%.[5] The relative contribution of these two components is governed by mixture parameter ϕ, inferred from the data.[6] We think of ϕ the *potential of a property to be present in a kind* and γ is the *mean prevalence of the property among the kinds with the potential to have it*.[7] If this model is correct, the prevalences given by participants would then be distributed as: $P(d) = \phi \cdot \text{Beta}(d \mid \gamma, \xi) + (1-\phi) \cdot \delta_{d=0}$.

We performed Bayesian inference over this model, given the observed prevalence data, to examine how well the model's posterior predictive distribution reconstructs the prevalence prior data. We put uninformative priors over all the parameters, $\phi \sim \text{Uniform}(0,1)$, $\gamma \sim \text{Uniform}(0,1)$, $\xi \sim \text{Uniform}(0,50)$, and performed Bayesian inference separately for each property using the Metropolis-Hastings algorithm collecting 50,000 samples removing the first 25,000 iterations for burn-in.

Results

Estimates of the mixture parameter ϕ and the mean of the "has the potential" component γ for each property are shown in Figure 2a. We see significant diversity among our properties in both parameters, corresponding to priors over prevalence with dramatically different shapes (insets).

Again, we look to the posterior predictive distribution to validate the structured prior model. Using the model with its inferred parameters, we generate prevalence judgments for different properties and compare that to the empirical counts. We discretize the prevalence values of both the model and the data to 12 discrete bins: $\{[0-0.01), (0.01-0.05), (0.05-0.15), (0.15-$

[5]There are other ways to formulate the second component ("the kind doesn't have a causal mechanism that would give rise to the property") of the prior. It could reflect accidental causes of the property, in which case, the prevalence could be a distribution that allows for non-zero prevalence. While an interesting possibility, its full consideration is beyond the scope of this article.

[6]This is similar in spirit to Hurdle Models of epidemiological data, where the observed count of zeros is often substantially greater than one would expect from standard models, such as the Poisson (e.g. adverse events to vaccinesRose, Martin, Wannemuehler, & Plikaytis, 2006)

[7]We note that ϕ is not what other authors have described as *cue validity* (Beach, 1964; Khemlani, Leslie, & Glucksberg, 2012), or $P(K \mid F)$. ϕ is a mixture component in the prior distribution over prevalence: $P(F \mid K)$. Cue validity and prevalence are related via Bayes Rule': $P(K \mid F) \propto P(F \mid K) \cdot P(K)$.

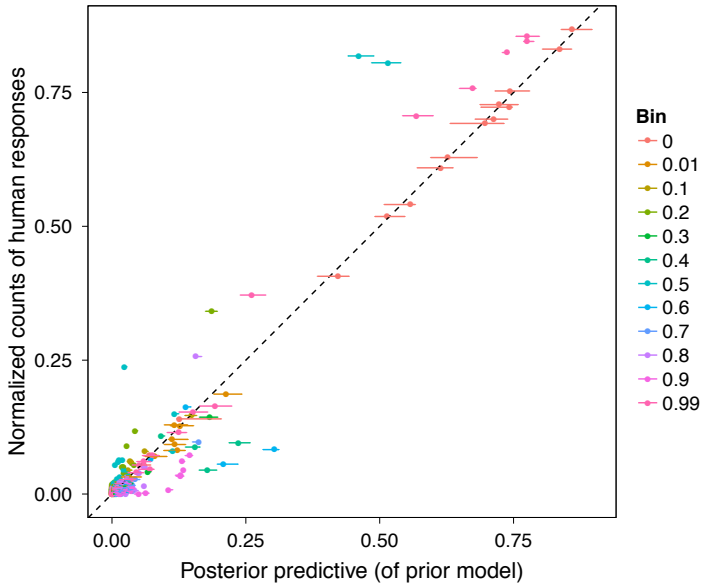

Figure 3: Posterior predictive distribution of the structured, statistical model thought to give rise to the human data in the prior elicitation task. The close alignment between model and data suggests the assumption of a structured prior is warranted.

$0.25), ..., (0.75-0.85), (0.85-0.95), (0.95-1]\}$. This statistical model reproduces the prior elicitation data very well ($r^2 = 0.94$), while a model that assumes just a single generative component fails ($r^2 = 0.14$). This is strong evidence in support of a structured prior.

The other test of this hypothesis is to re-examine the truth judgments from Expt. 1b using the pragmatics model with the inferred structured priors (as opposed to bootstrapping the raw empirical counts). We find the same correspondence to the empirical truth judgments data ($r^2 = 0.98$). This provides further evidence that the prior distribution over prevalence $P(x)$ is structured. The implication of this finding is that conceptual structure may indeed find its way into generic judgments, but via the prevalence prior, rather than directly in the semantics of the generic. We return to this idea in the General Discussion.

Empirical test 2: Interpreting novel generics

One of the most important roles for generic language is to provide learners information about new or poorly understood categories. This role depends on how unfamiliar generic sentences are interpreted (e.g. Gelman et al., 2002; Cimpian, Brandone, & Gelman, 2010). The pragmatic theory we present includes such a theory of generic comprehension: the listener model (Eq. 1) describes interpretation of a generic utterance—*Kind* HAS PROPERTY—without previously knowing the prevalence of the property within this kind. In our theory, the meaning is uncertain, but the

pressure to be informative operates over *a priori* beliefs about properties to produce an interpretation. Classic work in generalization suggests beliefs about the prevalence of properties differ by type of property, including relatively fine distinctions among properties that are all biological in nature (Nisbett, Krantz, Jepson, & Kunda, 1983). We leverage these diverse expectations, using properties that explore a wide range of *a priori* beliefs about prevalence.

Measuring *a priori* beliefs is tricky when the kinds are unknown. We cannot, as before, have participants fill out a table with rows corresponding to different animal kinds and columns corresponding to different properties: Nothing would distinguish the rows. Instead, we leverage the latent structure uncovered in our extended model analysis of Expt. 1 and decompose prevalence priors into 2 components: the property's potential to be present in a kind and the mean prevalence when present.

We use this novel method for measuring *a priori* beliefs about the prevalence of these properties for unfamiliar kinds (Expt. 2a). We then test the predictions of the pragmatic listener model L_1 using these empirically derived priors against human *interpretations* of novel generic sentences (Expt. 2b). Finally, we explain a previously reported empirical asymmetry between truth conditions and interpretations by comparing the speaker S_2 and listener L_1 models in the same experimental context (Expt. 2c).

Experiment 2a: Prevalence priors for unfamiliar kinds

Method

Participants

We recruited 40 participants over MTurk. All participants were native English speakers. The experiment took about 5-7 minutes and participants were compensated $0.75.

Procedure and materials

We constructed forty different properties to explore a wide range of *a priori* beliefs about prevalence. These items make up four categories of properties: body parts of a particular color (e.g. HAS GREEN FEATHERS), described vaguely (e.g. HAS SMALL WINGS), in accidental or disease states (e.g. HAS WET FUR, HAS SWOLLEN EARS), and without modification (e.g. HAS CLAWS). Because pilot testing revealed more variability for items in the accidental category relative to the other types of properties, we used twice as many exemplars of accidental properties, yielding a more thorough test of the quantitative predictive power of the L_1 interpretation model. We used 8 exemplars of each of the three non-accidental properties ("parts", "colored parts", "vague parts"),

and 16 exemplars of accidental properties, building on a stimulus set from Cimpian, Brandone, and Gelman (2010). All materials are shown in Table 3 in the Appendix.

In the task, participants were introduced to a "data-collection robot" that was tasked with learning about properties of animals. Participants were told the robot randomly sampled an animal to ask the participant about (e.g. The robot says: "We recently discovered animals called feps."). We then used a two-stage elicitation procedure, aimed to measure the two components of the structured prior model: (1) the potential of the property to be present in a kind and (2) the expected prevalence when present. To get at (1), the robot asked how likely it was that "there was *a* fep with PROPERTY" (potential to be present), to which participants reported on a scale from "unlikely" to "likely". For example, it is very likely that there is a fep that is female, less likely that there is a fep that has wings, and even less likely that there is a fep that has purple wings. To get at (2), the robot then asked, "Suppose there is a fep that has wings. What percentage of feps do you think have wings?" (expected prevalence when present). Participants completed a practice trial to make sure they understood the meanings of these two questions.

Data analysis and results

We used the same structured, statistical model for the prior data from Expt. 1. The only difference from Expt. 1a. is that our experimental data comes from inquiring about the parameters of the priors directly, as opposed to asking about particular samples from the prior (i.e. particular kinds) as was done in Expt. 1a. We assume these two measurements follow Beta distributions ($d_{potential} \sim \text{Beta}(\gamma_1, \xi_1)$; $d_{expected} \sim \text{Beta}(\gamma_2, \xi_2)$), and construct single prevalence distributions, $P(x)$, by sampling from the posterior predictive distribution of prevalence as we did before: $P(x) = \int [\phi \cdot \text{Beta}(x \mid \gamma_2, \xi_2) + (1-\phi) \cdot \delta_{x=0}] \cdot \text{Beta}(\phi \mid \gamma_1, \xi_1) d\phi$. We used the same uninformative priors over parameters ϕ, γ_i, ξ_i as in Expt. 1a.

Figure 4a shows a summary of the elicited priors, in terms of the diversity of $d_{potential}$ and $d_{expected}$. Biological properties are expected to be *a priori* more prevalent within a kind when present than accidental properties, with additional fine-grained differences within biological and accidental properties. Like the priors elicited using familiar categories, these priors elicited using unfamiliar categories have diverse shapes (see insets). Biological properties ("biological", "vague", and "color" body parts) have prevalence distributions that are bimodal with peaks at 0% and near-100% prevalence. Interpretations of generics about these properties (L_1 model, Eq. 1) update these distributions to concave posteriors peaked at 100% (Figure 4a; red, blue and green insets); the model predicts these novel generics will be interpreted as implying the property is widespread in the category. By contrast, accidental properties (both "rare" and "common") follow unimodal prior distributions and update to convex posterior distributions, predicting weaker and

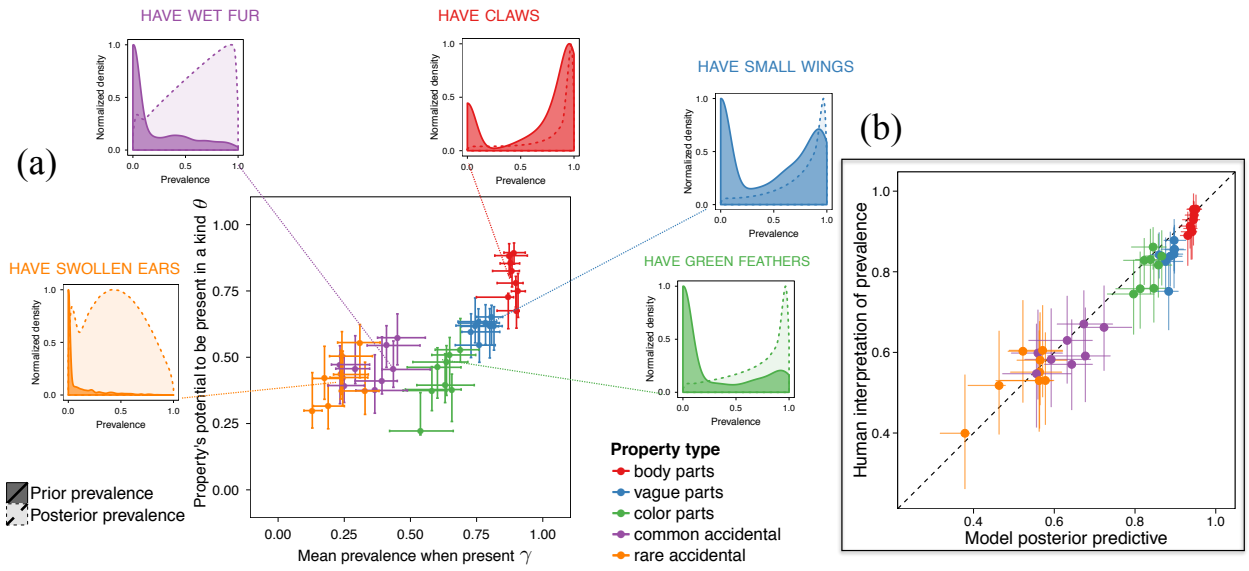

Figure 4: Understanding novel generics. (a) Prevalence prior distributions empirically elicited for 40 animal properties. Parameters of the structured statistical model—ϕ and γ—reveal quantitative differences in beliefs about the prevalence of conceptually different types of properties (scatterplot). Inset plots show differences in shapes between biological properties (red, green, blue; bimodal) and accidental properties (orange, purple; unimodal). These differences in the prior (darker shade) give rise to the variability of L_1 interpretations of generic utterances (lighter shade). (b) Human interpretation of prevalence upon hearing a generic compared with the L_1 model posterior predictive. Participants and the model interpret generics differently for different property types: Generics of biological properties (red, blue, green) have strong interpretations while generics of accidental properties (purple, orange) are weaker. Error bars denote Bayesian 95% credible intervals.

more variable interpretations of novel generics for these properties.

Experiment 2b: Interpretations of novel generics

Our model of generic interpretation, the pragmatic listener model L_1 (Eq. 1), predicts that the interpretations of generics in terms of prevalence should vary as a function of the prevalence prior. Here, we test the degree to which the predictions based on the empirically elicited prevalence priors for 40 items (from Expt. 2a) match human judgments of how the widespread the property is upon hearing a generic.

Method

Participants

We recruited 40 participants over MTurk to determine how widespread different properties are believed to be upon hearing a novel generic. The experimental design is very similar to Cimpian, Brandone, and Gelman (2010), and we chose to have a sample size at least twice as large as the original study (original n=15). All participants were native English speakers. The experiment took about 5 minutes and participants were compensated $0.60.

Procedure and materials

In order to get participants motivated to reason about novel kinds, they were told they were the resident zoologist of a team of scientists on a recently discovered island with many unknown animals; their task was to provide their expert opinion on questions about these animals. Participants were supplied with the generic (e.g., "Feps have yellow fur.") and asked to judge prevalence: "What percentage of feps do you think have yellow fur?". Participants completed in randomized order 25 trials: 5 for each of the biological properties and 10 for the accidental (described in Expt. 2a). The experiment in full can be viewed at `http://stanford.edu/~mtessler/generics/experiments/asymmetry/asymmetry-2.html`.

Analysis and results

The pragmatic listener L_1 model provides posterior beliefs about prevalence, given prior beliefs and a generic utterance. This model has one parameter governing the optimality of the hypothetical speaker S_1 in Eq. 2. We put the same uninformative prior over this parameter as previously: $\lambda_1 \sim \text{Uniform}(0, 20)$. We learned about the parameter's *a posteriori* credible values by running 3 MCMC chains of 100,000 samples (removing 50,000 for burn-in) using the Metropolis-Hastings algorithm. The MAP and 95% credible interval for λ_1 are 14.8 [6.4, 19.9].

We look at the posterior predictive distribution of L_1, integrating out the model parameter. We first explore two important trends predicted by the pragmatic listener model. In Figure 5 (solid lines) we see the implied prevalence judgments are predicted (at the property class level) to vary linearly with the *a proiri* expected prevalence. A mixed-effects linear model with random by-participant effects of intercept and slope indeed reveals the more prevalent a property is expected to be *a priori*, the stronger the implications of a generic statement ($\beta = 0.57; SE = 0.08; t(39) = 7.12; p < 0.001$). The prevalence implied by a generic is also predicted to be greater than the *a proiri* expected prevalence (i.e., greater than the prevalence expected among the kinds with the potential to have the property). A mixed-effects linear model with random by-participant effects

of intercept and random by-item effects of intercept and condition reveals implied prevalence after hearing a generic is significantly greater than the *a priori* prevalence ($\beta = 0.17; SE = 0.018; t(39) = 9.7; d = 0.64; p < 0.001$). As for the quantitative accuracy of the model, on a by-item level, the pragmatic listener model predictions closely align with the human judgments of prevalence for novel generics ($r^2(40) = 0.94$, MSE=0.002). Human participants and our model display the same sensitivity of generic interpretation to details of the property (Figure 4b). We now have strong support for both of the major predictive components of our model: generic endorsement, modeled as a speaker S_2, and generic interpretation, modeled as a listener L_1.

Experiment 2c: The asymmetry between truth conditions and interpretations

There is a surprising décolage between the truth conditions and interpretations of generic language: Interpretations are often strong while truth conditions are flexible. Cimpian, Brandone, and Gelman (2010) found that upon reading a generic (e.g. *Glippets have yellow fur.*), participants infer (in an *implied prevalence* task) that the property is widespread (e.g. almost all glippets have yellow fur). By contrast, participants endorse generics (in a *truth conditions* task) for a wide range of prevalence levels (e.g. even when "30% of glippets have yellow fur."), thus showing an asymmetry between truth conditions and implied prevalence. However, this mismatch is not found for the behavior of quantified statements involving "all" or "most," and is significantly reduced for generics of accidental properties (e.g. *Glippets have wet fur.*).

Below we replicate the basic asymmetry findings of Cimpian, Brandone, and Gelman (2010) and reveal even more variability in the mismatch between *truth conditions* and *implied prevalence* using the expanded stimulus set from Expt. 2a. In addition, we now test both our models (generic endorsement [speaker S_2] and generic interpretation [listener L_1]) in the same experimental paradigm.

Method

We re-analyze the data from Expt. 2b as the *implied prevalence* data. The following paradigm is to measure the corresponding *truth conditions*.

Participants

We recruited 40 participants over MTurk. All participants were native English speakers. None of the participants completed Expt. 2b (interpretations of novel generics). The experiment took about

5 minutes and participants were compensated $0.60.

Procedure and materials

The cover story and materials were the same as in Expt. 2b. On each trial, participants were given a statement about a property's prevalence within a novel kind (e.g. *50% of feps have yellow fur.*). Participants were then asked whether or not they agreed or disagreed with the corresponding generic sentence (e.g. *Feps have yellow fur.*). Prevalence varied between 10, 30, 50, 70, and 90%.

The experiment consisted of 25 trials: 5 trials for each of 5 types of properties measured in Expt. 2a (part, color part, vague part, common accidental, rare accidental). Each prevalence level appeared once for each property type (5 prevalence levels x 5 property types).

Analysis and results

For both behavioral data and model predictions (Eq. 4) we computed the average prevalence that led to an assenting judgment (the *average prevalence score*), for each property type and participant, following the procedure used by Cimpian, Brandone, and Gelman (2010). For example, if a participant agreed with the generic whenever the prevalence was 70% or 90% and disagreed at the other prevalence levels, that participant received an *average prevalence score* of 80%.

For our pair of models, there are two parameters (the two speaker optimality parameters). We infer them using the same Bayesian data analytic approach as before. The MAP and 95% HPD intervals for λ_1 is 19.5 [10.5, 19.9] and λ_2 is 0.4 [0.34, 0.49]. We then subjected the generic endorsement model to the same procedure as the human data. The speaker model S_2 returns a posterior probability of producing the generic, for each level of prevalence. We sample a response (*agree* / *disagree*) from this posterior distribution for each prevalence level, simulating a single subject's data. As with the human data, we took the trials where the model agreed with the generic, and took the mean of the prevalence levels corresponding to those trials, giving us the average prevalence at which the model assented to the generic. We repeated this for each type of property 40 times to simulate a sample of 40 participants. We repeated this procedure 1000 times to bootstrap 95% confidence intervals.

The generic endorsement model (speaker S_2) predicted that *average truth conditions* should not vary appreciably across the different types of properties, consistent with the fact that generics are acceptable for broad range of prevalence levels for all property types. A similar absence of a gradient was observed in the human data ($\beta = 2.82; SE = 4.02; t(39) = 0.70; p = 0.49$; Figure 5, dotted lines). Interpretations of generic utterances are stronger than their average truth conditions for the biological properties but not for the accidental properties (Figure 5) with both human data, replicating Cimpian, Brandone, and Gelman (2010), and the model; the extent of the difference is

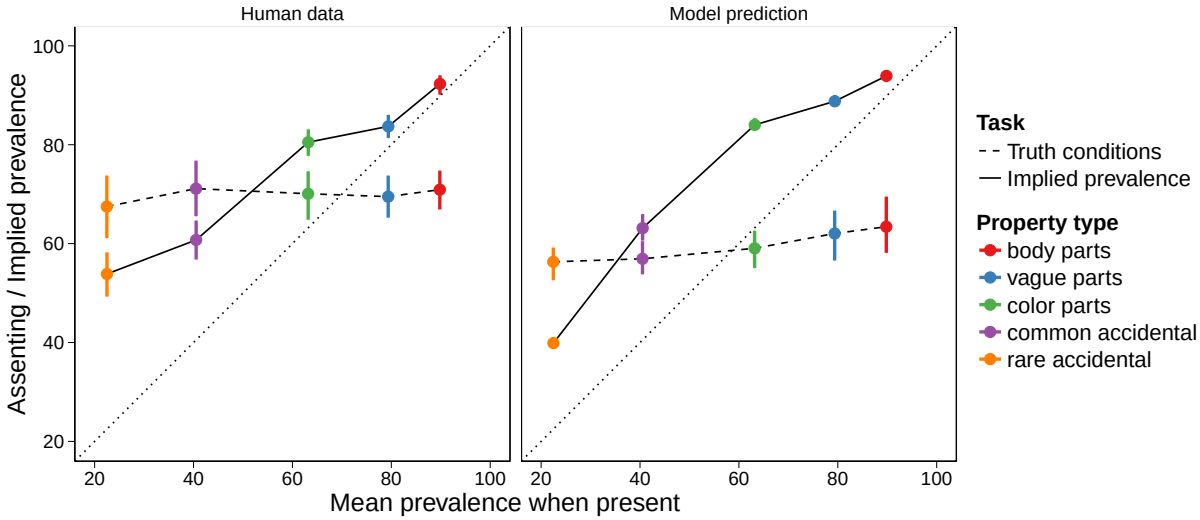

Figure 5: The asymmetry between truth conditions and interpretations. Human judgments and model predictions of prevalence implied by novel generic utterances (implied prevalence task; solid line) and average prevalence that leads to an acceptable generic utterance (truth conditions task; dotted line) as it relates to the *a priori* mean prevalence when present γ. Expectations of prevalence are higher after hearing a generic than before hearing it (solid line compared to $y = x$ line; both for human data and model). Generic statements about biological properties, imply that the property is widespread in the category, for both human participants and the model (solid line: red, blue and green). Generics about accidental properties do not result in such a high implied prevalence (solid line: purple and orange). While the implications of generic utterances are highly variable across the different types of properties, the average prevalence that leads to an acceptable generic does not vary, for participants nor the model.

governed by prior property knowledge (mean prevalence when present γ, from Expt. 2a). The listener and speaker pair of models predicts human endorsements and interpretations of novel generic utterances well ($r^2(10) = 0.87$, MSE = 0.008). Thus, our model predicts that the asymmetry between truth conditions and implied prevalence should hold, but only for properties with the most extreme prior beliefs.

Prevalence is a predictive probability

So far, we have shown that property prevalence is sufficient to formalize the semantics of generic statements as an underspecified scalar denotation. But what is property prevalence? If generic language is truly conveying generalizations, it would be useful for it to reflect expectations, not just

the current statistics in the world. The current frequency of a property is often a good indicator of future frequency, yet statistics can be distorted by spurious events. The causal history of a property may be more or less important for implying the property will be present in future situations. Does generic language communicate prevalence in terms of past frequency or future expectations?

To answer this, we adopt an experimental paradigm used by Gelman and Bloom (2007) to show that generic language is sensitive to theory-based considerations. In the original paradigm, participants are told a story about a novel creature (e.g. *dobles*) and a property of that kind (e.g., *having claws*). Participants are then either told that the creature was born with the property or that it acquired the property through extrinsic means (e.g., by finding claws and putting them on). Then, participants are told about an event that either causes the property to disappear (e.g., they drank a chemical and their claws fell off) or that leaves the property intact, and are asked whether or not the generic (e.g. *Dobles have claws*) applies. The original finding was that adult judgments were sensitive to the origins of the property (i.e., born vs. acquired), and insensitive to the outcome of the event (i.e., property maintained vs. lost): Participants fully-endorsed the generic when it was inborn, and rejected it when it was acquired, regardless of the current prevalence of the property.

In Experiment 3a, we use the same basic paradigm to measure *predictive prevalence*: participants' expectations about future instances of the kind. We explore the predictions of our truth judgments model, assuming that *predictive prevalence* is what is being communicated. In Experiment 3b, we use a truth judgment task similar to Gelman and Bloom (2007) and compare participants' judgments to the model's predicted endorsements.

Experiment 3a: Predictive prevalence elicitation

The design of this experiment is based on a study reported in Gelman and Bloom (2007) with some slight modification.

Method

Participants

We recruited 80 participants over MTurk. The experiment took about 3 minutes and participants were compensated $0.35.

Procedure and materials

On each trial, participants read a vignette about a novel creature. For instance,

> These are dobles. [picture of 10 dobles with claws] Here is how they grew. They grew up with claws. First they were born, then they got bigger, then they were full size. [picture of a doble with claws, getting bigger and bigger; in some vignettes, the animal was first shown hatching out of an egg with the relevant property already visible] Then one day they drank a bad chemical. They got very sick and this is how they looked. [picture of 10 dobles without claws]

The trial proceeded by participants reading the text, and clicking a button to continue to the next part of the story (at which time, the images changed according to the example above).

Participants saw 4 trials: 2 in which the creatures are born with the property (**intrinsic origins**), and 2 in which the creatures are shown discovering and acquiring the property (**extrinsic origins** e.g., painting themselves brown). This was crossed with either the creatures drinking a "bad chemical" and losing the feature, or drinking a "yummy drink" and maintaining the feature. The outcome of this event determined the final presentation of images that the participant saw (e.g., either 10 dobles with claws or 10 without).

While this final screen was present, we measured *predictive prevalence* by telling participants: "A new doble was born today. When it becomes full grown, how likely is it that it would have claws?" Participants responded using sliders ranging from "very unlikely" to "very likely".

We used 2 different types of properties: colors (e.g. *Lorches are green.*) and body parts (e.g. *Dobles have claws*). For each type of property, there were approximately 8 different exemplars (different colors or different body parts for different creatures). The creatures were either birds, bugs, or fish, with randomly sampled physical dimensions (e.g., sizes of body or tail). The experiment in full can be viewed at `http://stanford.edu/~mtessler/generics/experiments/predictive/predictive-elicitation-1-elicitation.html`.

Results and truth judgment predictions

The average predicted prevalences for the 4 experimental conditions are shown in Table 1. We observe a main effect of origins, such that when participants read that the creatures had the property from birth, future creatures are much more likely to have the property as compared to when the property is acquired. We see that, in our paradigm, participants are also sensitive to the outcome of the event. When participants observe a creature who loses the property by drinking a chemical, they report future members of the category are less likely to have the property. This inference may be driven by inferences about the property (e.g., that the property could be an unstable property, if you can lose it simply by drinking something) or by inferences about the event (e.g., participants may believe this "chemical drinking" event is a relatively normal event, and thus it could happen in the future).

Origins	Event Outcome	Collapsed Mean [95% CI]	Color [95% CI]	Body Part [95% CI]
Extrinsic	Lost	0.15 [0.10, 0.21]	0.13 [0.07, 0.21]	0.17 [0.09, 0.26]
Extrinsic	Maintained	0.32 [0.24, 0.39]	0.24 [0.16, 0.35]	0.38 [0.27, 0.50]
Intrinsic	Lost	0.69 [0.62, 0.76]	0.73 [0.63, 0.83]	0.66 [0.56, 0.75]
Intrinsic	Maintained	0.95 [0.94, 0.97]	0.96 [0.94, 0.98]	0.94 [0.91, 0.97]

Table 1: Predicted prevalence for the four experimental conditions of Expt. 3a. Right two columns show the summaries broken down by the two types of properties used in the experiment.

We use these predicted probabilities as the prevalence x that the speaker model is trying to communicate: $S_2(u \mid x)$, and examine the model's predicted truth judgments. We explore the model's predictions for each origin and event outcome, as well as when the data is split by property type (color vs. body parts). For priors $P(x)$, we use the body part and color priors elicited in Expt. 2a. We see that the model predictions track closely the predicted prevalence (Figure 6a, top, compare with predicted prevalence in Table 1). This is because both color and body part priors are relatively broad, and hence when the property is (predicted to be) more prevalent, the generic has a higher probability of applying (see schematic predictions from Figure 1 "have wings" for comparison). We also see that the model predicts a subtle by-item difference, such that the influence of the event outcome (lost or maintained) on generic endorsement is predicted to be stronger for body parts than for color terms (Figure 6a, bottom). This prediction is mostly due to the predicted prevalence for the conflict conditions (intrinsic-lost and extrinsic-maintained) being subtly different (Table 1, right-most columns).

Our model thus makes two novel predictions for generic endorsement in the paradigm by Gelman and Bloom (2007). We predict that in addition to the main effect of origins, we should see a second main effect of event outcome. Second, we predict that this effect should be slightly stronger in the case of color properties than in the case of body part properties.

Experiment 3b: Truth judgment task

In this experiment, we test the predictions of our model using a truth judgment measure in the same paradigm.

Participants

We recruited 80 participants over MTurk. The experiment took about 3 minutes and participants were compensated $0.35. None of the participants completed Experiment 3a as well.

Procedure and materials

The procedure and materials are exactly the same as in Expt. 3a, with the exception of the dependent measure. After reading each vignette, participants were asked: "Do you agree or disagree that: GENERIC STATEMENT (e.g. Dobles have claws)". Participants responded by choosing one of two radio buttons corresponding to agree or disagree.

Results and discussion

Our pragmatics model with the elicited predicted prevalence from Expt. 3a made two novel predictions for this experiment: (1) in addition to a main effect of origins, we would find a main effect of event outcome; (2) this effect would be stronger for body part properties than for color properties. As predicted, we found two main effects (Figure 6b, top). The main effect of property origins replicated ($\beta = 3.6, SE = 0.57, z = 6.5, p < 1e10$): participants were more likely to endorse the generic when it was about a property that the creature was born with. In addition, we find a second main effect of event outcome ($\beta = 2.69, SE = 0.56, z = 4.8, p < 1e5$): participants were more likely to endorse the generic when the property was maintained than when it was lost.[8]

When we break down the results by item, we see that this effect is stronger for body part properties than for color properties (Figure 6b, bottom). The endorsement of a generic for color properties (e.g. *Lorches are green*) seems to be less sensitive to the outcome of the event (i.e. Lorches losing their color as a result of drinking a chemical). This may be due to participants' intuitive theories of properties and their stability (skin color is more stable than body parts like feathers). Indeed this difference is apparent in the predictive prevalence task (Expt. 3a). For the 8 data points of generic endorsement based on origins, outcome, and property type, our model's predictions match the data well ($r^2(8) = 0.96$). We, thus, elaborate our theory: The semantics of generics can be understood as a threshold on property prevalence, and this prevalence is a speaker's subjective belief about what is likely to be the case in the future.

General discussion

Generic language is the simple and ubiquitous way by which generalizations are conveyed between people. Yet the dramatic flexibility of generic language has confounded psychologists, linguists and philosophers who have tried to articulate what exactly generic statements mean. We evaluated

[8]The fact that we find a second main effect of event outcome in addition to origins, whereas the original only found a main effect of origins, makes it worth noting the differences between our paradigm and the original study by Gelman and Bloom. In the original study, the first sentence of each vignette used the possessive "my": "These are my dobles.". At the end of each vignette, the original study had participants judge two statements in counterbalanced order: "Do my dobles have claws?" and "Do dobles have claws?" Finally, the original sample size was 14; ours was 80.

(a) S_2 model predictions

(b) Human endorsement of generic statements

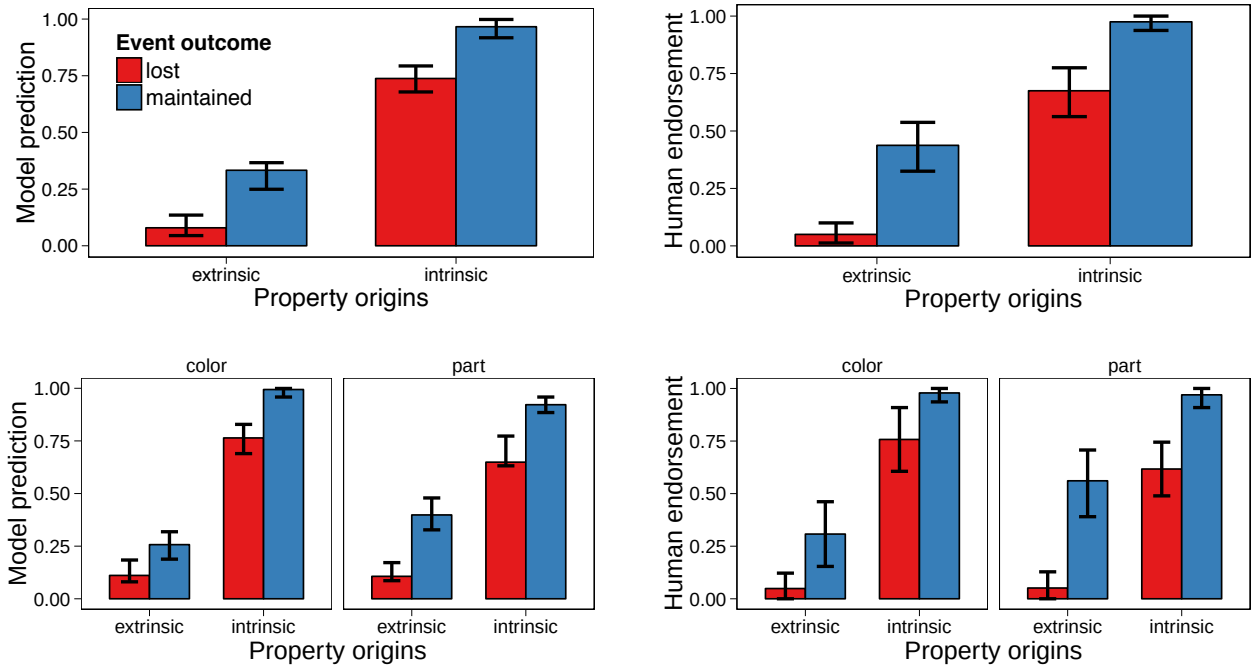

Figure 6: Prevalence is a predictive probability. (a) Truth judgment model predictions given the predicted prevalence elicited in Expt. 3a. (b) Average endorsement of the generic statement in Expt. 3b (replication of Gelman and Bloom, 2007). Bottom row shows data and predictions broken down by property type.

a theory of generic language derived from general principles of language understanding using a simple, but uncertain, basic meaning—a threshold on property prevalence. Our formal model is a minimal extension of the RSA theory of language understanding, together with an underspecified threshold semantics. The model was able to explain two major puzzles of generics: their extreme flexibility in truth conditions and the contrastingly strong interpretation of many novel generics. Both of these phenomena were revealed to depend in systematic ways on prior knowledge about properties. This prior knowledge was revealed through Bayesian model analysis to be structured, providing a promising bridge to conceptual accounts of generic language. To understand the nature of the underlying prevalence scale, we showed that generic language is about speakers' expectations of future prevalence, and not necessarily what the current state of the world is like. Across all experiments, the formal model predicted the quantitative details of participants' judgments with high accuracy.

There have been numerous demonstrations arguing that statistics (e.g., property prevalence) are insufficient to explain generic meaning (Gelman et al., 2002; Gelman & Bloom, 2007; Cimpian, Brandone, & Gelman, 2010; Cimpian, Gelman, & Brandone, 2010; Khemlani et al., 2012; Prasada et al., 2013). In these experiments, the prevalence considered is only the prevalence of the property for the target category (e.g., the percentage of birds that lay eggs; Khemlani et al., 2012; Prasada et al., 2013), what we have referred to as *within-kind prevalence*. Indeed, this simple statistic also fails to explain our data (Figure 2b, right). Our formal model of pragmatics, by contrast, considers not only within-kind prevalence, but a listener's prior beliefs about prevalence across kinds in order to arrive at a meaning for a generic utterance. By establishing the validity of a semantics based on prevalence alone, we provide a formalism to learn about categories from generic statements. Further, since prevalence is a probability, our model can take information conveyed with a generic and be naturally extended to make predictions about entities in the world or support explanations of events or behavior.

The comparison class

In this paper, we proposed a model for understanding generic language that relies upon interlocutors' belief distribution over the statistics of the property to accurately arrive at generic interpretation identical to that of human participants. The prior belief distribution for a property (e.g., lays eggs), is a distribution over kinds (and their associated prevalence of the property). These other kinds form a *comparison class* against which the target kind is evaluated. $P(x)$ is always relative to a category of categories, a comparison class K: $P_K(x)$.

Throughout our experiments in this work, we have focused on sentences about animals. In addition to being the main focus of past theoretical and empirical work, focusing on animals is

methodologically convenient as the comparison class for generics about animals is quite naturally *animals*. When we look beyond generics about animals, deciding what goes into a comparison class becomes less clear. There are some hints that the comparison class can be derived with respect to the property (Keil, 1979), but may involve pragmatic reasoning as well. For example, the statement "iPhones are useful" could be in comparison to other forms of technology (like a desktop computer), while "iPhones are heavy" could really only be informative relative to other handheld devices.

The incorporation of a comparison class into the study of generic language might help elucidate other puzzles concerning generics. Recent work in philosophy and linguistics, for instance, suggest generic language is context-sensitive (Nickel, 2008; Sterken, 2015). Nickel (2008) argues that *Dobermans have floppy ears* may be true in the context of a discussion of evolutionary biology but that *Dobermans have pointy ears* is true in a discussion of dog breeding. Our theory provides a hint from where to begin to understand this context sensitivity: the comparison class. Different conversational contexts could bring to mind different comparison classes, in a way analogous to the context-sensitivity of gradable adjectives (e.g., *tall*). Hearing that "Abigail is tall" means different things is Abigail is 20 years old or if she is 4. Future work will be needed to explore whether a pragmatic inference approach is also relevant to establishing the comparison class, and what background knowledge about properties, categories, and context is relevant.

Generic identification

Throughout this paper we treated the bare plural construction as a generic utterance with a threshold semantics: $[\![\text{K F}]\!](x,\theta) = x > \theta$. The bare plural construction can also indicate a specific plural predication. For example, "Dogs are on my lawn" picks out a specific group of dogs, while "Dogs have fur" does not (Carlson, 1977). The problem of *identifying* a generic meaning from a bare plural construction is itself a challenging problem because generic meaning can be signaled using a diverse array of morphosyntactic cues.

Declerck (1991) suggests that generic and non-generic bare plurals can be treated in the same way, and that pragmatic considerations alone may resolve interpretative differences. Indeed it does appear that knowledge of the properties under discussion (e.g., the state of being on a front lawn; the state of having fur) could facilitate the generic identification process. Other pragmatic factors, like knowledge of the identity of the speaker (e.g., a teacher vs. a veterinarian), can also disambiguate generic and non-generic meaning (Cimpian & Markman, 2008). Recent work suggests that utterances that fail to refer to specific entities or events could pragmatically imply generic meaning (Crone & Frank, 2016). Incorporating these insights about generic identification into an information-theoretic, communicative perspective is a natural extension of this work.

Implications for conceptual structure

Previous psychological and philosophical work on generics has looked beyond prevalence and focused on conceptual distinctions and relations (Gelman, 2003; Prasada et al., 2013; Leslie, 2007, 2008). Prasada et al. has argued for a distinction between *characteristic* properties (e.g. *Diapers are absorbent.*) and *statistical* properties (e.g. *Diapers are white.*). Leslie suggests information that is striking (e.g. *Tigers eat people.*) is useful and thus permitted to be a generic. Gelman outlines how generics tend to express *essential* qualities that are relatively timeless and enduring. Where in the prevalence-based semantics could such conceptual distinctions come into play?

Our approach makes the strong claim that beliefs about predicted prevalence are the connective tissue between conceptual knowledge and generic language. That is, the effect of conceptually meaningful differences on generic language is predicted to be mediated by differences in corresponding prevalence distributions. It is important to note that our approach is based on *subjective* probability, and not mere frequency. Indeed, we elucidated in Expt. 3 that using participants' *predictions* of probability in our formal model perfectly track generic endorsement, when the present frequency would make the wrong prediction.

The focus on subjective, predictive probability casts new light on puzzles surrounding accidentally-true situations. An example is the statement "Supreme Court Justices have even social security numbers", which is predicted by linguists to be rejected even if every single Supreme Court Justice has an even social security number (Cohen, 1999). Our explanation is that abstract intuitive theories lead us to reject observed frequencies in forming our subjective probabilities. That is, because we may believe selection for the Supreme Court is not influenced by one's social security number, we would assign roughly 50% subjective probability to the *next* justice having an even number. Thus, we would still reject the generic *Supreme Court Justices have even social security numbers*, because the predictive prevalence would not be any different for Supreme Court Justices than any other profession.[9]

Turning back to conceptual relations and structure, it is natural to ask when subjective probabilities might reflect conceptual knowledge? We found that empirical prevalence distributions are structured in a way that reflects intuitions about causal mechanisms underlying different properties; the differences in shape of these distributions in turn led to variable endorsements and interpretations of generic sentences. It is plausible that richer conceptual knowledge also influences these distributions, such as higher-order conceptual knowledge about the nature of properties and categories (Gelman, 2003; Keil, 1992). Indeed, it has been argued that conceptual structure in general,

[9]Our perspective makes the intriguing prediction that if we learned much more surprising information, we might be compelled to revise our theory and then accept the generic. For instance, if every justice in history had *prime numbered* social security numbers (a more suspicious coincidence), one might appeal to a conspiracy, which *would* have predictive consequences.

including higher-order abstractions, can be captured by probabilisitic causal models and their generalizations (Pearl, 1988; Gopnik, 2003; Goodman, Tenenbaum, & Gerstenberg, 2015). Future work will be needed to explore whether probabilistic representations of conceptual knowledge can capture the relations identified in other accounts of generics (such as characteristic, essential, and striking properties), and whether the effect of these relations can then be adequately captured via their impact on subjective prevalence.

Conclusion

It might seem paradoxical that a part of language that is so common in communication and central to learning should be vague. Shouldn't speakers and teachers want to express their ideas as crisply as possible? To the contrary, underspecification can be efficient, given that context can be used to resolve the uncertainty (Piantadosi, Tily, & Gibson, 2012). In our work, context takes the form of a listener and speaker's shared beliefs about the property in question. By leveraging this common ground, generics provide a powerful way to communicate and learn generalizations about categories, which would otherwise be difficult or costly information to learn through direct experience.

The dark side of this flexibility is the potential for miscommunication or deceit: A speaker might assert a generic utterance that he himself would not accept, conveying a too-strong generalization to a naïve listener. Our model predicts this potential particularly for properties which, when present, are widespread in a category—we showed that biological properties are believed to have this distribution, but many properties of social categories may as well (Cimpian & Markman, 2011; Cimpian, Mu, & Erickson, 2012; Rhodes et al., 2012). Disagreements are also predicted when interlocutors fail to share background assumptions: differences in the within-kind prevalence, the prior distributions on prevalence, or the comparison class. For example, there is considerable disagreement as to whether or not "Humans cause global warming". Our theory predicts this disagreement may be the result of differences in the estimated *causal power* of humans influencing global warming as well as the causal power of *other forces* (e.g., plate tectonics) on climate change. This is a promising area for future research.

Categories are inherently unobservable. You cannot see the category DOG, only some number of instances of it. Yet we easily talk about these abstractions, conveying hard-won generalizations to each other and down through generations. The theory presented here gives one explanation of how we do so, providing a computational perspective on how category generalizations are conveyed and how beliefs play a central role in understanding language.

References

Anderson, J. R. (1991). Chapter 1: Introduction. In *Adaptive character of thought* (pp. 1–38).

Baker, C. L., Saxe, R., & Tenenbaum, J. B. (2009). Action understanding as inverse planning. *Cognition*, *113*, 329–349. Retrieved from http://dx.doi.org/10.1016/j.cognition.2009.07.005 doi: 10.1016/j.cognition.2009.07.005

Beach, L. R. (1964). Cue probabilism and inference behavior. *Psychological Monographs: General and Applied*, *78*(5), 1–20. doi: 10.1017/CBO9781107415324.004

Behrens, L. (2005). Genericity from a cross-linguistic perspective. *Linguistics*, *43*(2), 275–344.

Berger, J. (1988). *Statistical decision theory and bayesian analysis.* New York: Springer-Verlag.

Brandone, A. C., Gelman, S. A., & Hedglen, J. (2014). Children's Developing Intuitions About the Truth Conditions and Implications of Novel Generics Versus Quantified Statements. *Cognitive science*, 1–28.

Carlson, G. N. (1977). *Reference to kinds in english*. Unpublished doctoral dissertation, University of Massachusetts, Amherst.

Carlson, G. N. (1995). Truth conditions of generic sentences: Two contrasting views. In G. N. Carlson & F. J. Pelletier (Eds.), *The generic book* (p. 224-38). Chicago: University of Chicago Press.

Carlson, G. N., & Pelletier, F. J. (1995). *The generic book.* Chicago, IL: Chicago University Press.

Cimpian, A. (2010). The impact of generic language about ability on children's achievement motivation. *Developmental psychology*, *46*(5), 1333–1340. doi: 10.1037/a0019665

Cimpian, A., Brandone, A. C., & Gelman, S. A. (2010). Generic statements require little evidence for acceptance but have powerful implications. *Cognitive science*, *34*(8), 1452–1482.

Cimpian, A., Gelman, S. A., & Brandone, A. C. (2010). Theory-based considerations influence the interpretation of generic sentences. *Language and Cognitive Processes*, *25*(2), 261–276. doi: 10.1080/01690960903025227.Theory-based

Cimpian, A., & Markman, E. M. (2008). Preschool children's use of cues to generic meaning. *Cognition*, *107*, 19–53. doi: 10.1016/j.cognition.2007.07.008

Cimpian, A., & Markman, E. M. (2011). The Generic/Nongeneric Distinction Influences How Children Interpret New Information About Social Others. *Child Development*, *82*(2), 471–492. doi: 10.1111/j.1467-8624.2010.01525.x

Cimpian, A., Mu, Y., & Erickson, L. C. (2012). Who Is Good at This Game? Linking an Activity to a Social Category Undermines Children's Achievement. *Psychological Science*, *23*(5), 533–541. doi: 10.1177/0956797611429803

Clark, H. H. (1996). *Using language*. Cambridge University Press.

Cohen, A. (1999). Generics, Frequency Adverbs, and Probability. *Linguistics and Philosophy*, *22*.

Cohen, A. (2004, October). Generics and Mental Representations. *Linguistics and Philosophy*, *27*(5), 529–556. Retrieved from http://link.springer.com/10.1023/B:LING.0000033851.25870.3e doi: 10.1023/B:LING.0000033851.25870.3e

Crone, P., & Frank, M. C. (2016). Inferring generic meaning from pragmatic reference failure. In *Proceedings of the 37th annual meeting of the cogntiive science society.*

Declerck, R. (1991). The origins of genericity. *Linguistics*, *29*, 79–102.

Degen, J., & Goodman, N. D. (2014). Lost your marbles? the puzzle of dependent measures in experimental pragmatics. In *Proceedings of the thirty-sixth annual conference of the Cognitive Science Society.*

Frank, M. C., & Goodman, N. D. (2012). Predicting pragmatic reasoning in language games. *Science*, *336*(6084).

Franke, M. (2009). *Signal to act: Game theory in pragmatics*. Unpublished doctoral dissertation, Universiteit van Amsterdam.

Franke, M., & Jäger, G. (2015). Probabilistic pragmatics, or why Bayes' rule is probably important for pragmatics. In *Zeitschrift für sprachwissenschaft* (pp. 3–44).

Gelman, S. A. (2003). *Essential child: Origins of essentialntialism in everyday thought.* Oxford University Press.

Gelman, S. A. (2004). Learning words for kinds: Generic noun phrases in acquisition. In *Weaving a lexicon* (p. 445-484). MIT Press.

Gelman, S. A., & Bloom, P. (2007). Developmental changes in the understanding of generics. *Cognition*, *105*(1), 166–183. doi: 10.1016/j.cognition.2006.09.009

Gelman, S. A., Goetz, P. J., Sarnecka, B. W., & Flukes, J. (2008). Generic Language in Parent-Child Conversations. *Language Learning and Development*, *4*(1), 1–31. doi: 10.1080/15475440701542625.Generic

Gelman, S. A., Star, J. R., & Flukes, J. E. (2002). Children's Use of Generics in Inductive Inferences. *Journal of Cognition and Development*, *3*(2), 179–199.

Gelman, S. A., Taylor, M. G., Nguyen, S. P., Leaper, C., & Bigler, R. S. (2004). Mother-child conversations about gender: Understanding the acquisition of essentialist beliefs. *Monographs of the Society for Research in Child Development*, *69*(1), vii, 116–127. doi: 10.1111/j.1540-5834.2004.06901001.x

Goodman, N. D., & Frank, M. C. (To appear). Pragmatic language interpretation as probabilistic inference. *Trends in Cognitive Sciences*.

Goodman, N. D., & Lassiter, D. (2015). Probabilistic semantics and pragmatics: Uncertainty in language and thought. In S. Lappin & C. Fox (Eds.), *The handbook of contemporary semantic*

theory, 2nd edition. Wiley-Blackwell.

Goodman, N. D., & Stuhlmüller, A. (2013). Knowledge and implicature: Modeling language understanding as social cognition. *Topics in Cognitive Science*.

Goodman, N. D., & Stuhlmüller, A. (2014). *The Design and Implementation of Probabilistic Programming Languages.* http://dippl.org. (Accessed: 2015-7-17)

Goodman, N. D., Tenenbaum, J. B., & Gerstenberg, T. (2015). Concepts in a probabilistic language of thought. In *The conceptual mind: New directions in the study of concepts.* MIT Press.

Gopnik, A. (2003). The theory theory as an alternative to the innateness hypothesis. *Chomsky and his critics*, 238–254.

Grice, H. P. (1975). Logic and conversation. In *Readings in language and mind.* Blackwell.

Griffiths, T. L., & Tenenbaum, J. B. (2005, December). Structure and strength in causal induction. *Cognitive psychology*, *51*(4), 334–84. doi: 10.1016/j.cogpsych.2005.05.004

Keil, F. C. (1979). *Semantic and conceptual development.* Harvard University Press.

Keil, F. C. (1992). *Concepts, kind, and cognitive development.* MIT Press.

Khemlani, S., Leslie, S.-J., & Glucksberg, S. (2012, July). Inferences about members of kinds: The generics hypothesis. *Language and Cognitive Processes*, *27*(6), 887–900. Retrieved from http://www.tandfonline.com/doi/abs/10.1080/01690965.2011.601900 doi: 10.1080/01690965.2011.601900

Lassiter, D., & Goodman, N. D. (2013). Context, scale structure, and statistics in the interpretation of positive-form adjectives. In *Semantics and Linguistic Theory (SALT) 23.*

Lassiter, D., & Goodman, N. D. (2015). Adjectival vagueness in a bayesian model of interpretation. *Synthese*.

Lee, M. D., & Wagenmakers, E. (2014). *Bayesian cognitive modeling: A practical course*. Cambridge: Cambridge University Press.

Leslie, S.-J. (2007). Generics and the Structure of the Mind. *Philosophical Perspectives*, *21*(1), 375–403.

Leslie, S.-J. (2008, July). Generics: Cognition and acquisition. *Philosophical Review*, *117*(1).

Leslie, S.-J., Cimpian, A., Meyer, M., & Freeland, E. (2015). Expectations of brilliance underlie gender distributions across academic disciplines. *Science*, *347*(6219), 262–265. doi: 10.1081/E-EWS

Levinson, S. (2000). *Presumptive meanings: The theory of generalized conversational implicature*. The MIT Press.

Markman, E. M. (1989). *Categorization and naming in children: Problems of induction*. MIT Press.

Marr, D. (1980). Chapter 1: The Philosophy and the Approach. In *Vision.*

Nickel, B. (2008). Generics and the ways of normality. *Linguistics and Philosophy*, *31*(6), 629–648. doi: 10.1007/s10988-008-9049-7

Nisbett, R. E., Krantz, D. H., Jepson, C., & Kunda, Z. (1983). The use of statistical heuristics in everyday inductive reasoning. *Psychological Review*, *90*(4), 339–363. doi: 10.1037/0033-295X.90.4.339

Pearl, J. (1988). *Probabilistic reasoning in intelligent systems: Networks of plausible reasoning*. Morgan Kaufmann Publishers, Los Altos.

Piantadosi, S. T., Tily, H., & Gibson, E. (2012). The communicative function of ambiguity in language. *Cognition*, *122*(3), 280–291. doi: 10.1016/j.cognition.2011.10.004

Prasada, S. (2000). Acquiring generic knowledge. *Trends in Cognitive Sciences*, *4*(2), 66–72. doi: 10.1016/S1364-6613(99)01429-1

Prasada, S., & Dillingham, E. M. (2006). Principled and statistical connections in common sense conception. *Cognition*, *99*(1), 73–112. doi: 10.1016/j.cognition.2005.01.003

Prasada, S., Hennefield, L., & Otap, D. (2012). Conceptual and Linguistic Representations of Kinds and Classes. *Cognitive Science*, *36*(7), 1224–1250. doi: 10.1111/j.1551-6709.2012.01254.x

Prasada, S., Khemlani, S., Leslie, S.-J., & Glucksberg, S. (2013, March). Conceptual distinctions amongst generics. *Cognition*, *126*(3), 405–22. doi: 10.1016/j.cognition.2012.11.010

Rhodes, M., Leslie, S.-J., & Tworek, C. M. (2012). Cultural transmission of social essentialism. *Proceedings of the National Academy of Sciences*, *109*(34), 13526–13531. doi: 10.1073/pnas.1208951109

Rose, C. E., Martin, S. W., Wannemuehler, K. A., & Plikaytis, B. D. (2006). On the use of zero-inflated and hurdle models for modeling vaccine adverse event count data. *Journal of biopharmaceutical statistics*, *16*(4), 463–481.

Sterken, R. K. (2015). Generics in Context. *Philosophers' Imprint*, *15*(i), 1–30.

Appendix A: The Rational Speech Act framework

The goal of this section is introduce the reader to the Rational Speech Act (RSA) framework. Goodman and Frank (To appear) provides another introduction and useful review of the framework.

The Rational Speech Act theory views language understanding as a special case of social cognition instantiated in a Bayesian model (Frank & Goodman, 2012; Goodman & Stuhlmüller, 2013; Goodman & Frank, To appear; see also, Franke, 2009; Franke & Jäger, 2015). In this view, an utterance is *interpreted* by a Bayesian listener by integrating her prior beliefs with the likelihood that a speaker would have generated the utterance. This generative model of the utterance is the listener's intuitive theory of speech *production*, i.e. her theory of a speaker. This intuitive theory of speech production views speakers as rational actors who behave in accord with Bayesian decision-theory (Berger, 1988). The rational speaker, thus, chooses utterances soft-max optimally on the basis of conveying information to a naive listener. The pragmatic listener, then, "inverts" her model of a speaker to arrive at a pragmatic interpretation of an utterance.

This model is posed at the computational or rational level (Marr, 1980; Anderson, 1991). It is not intended as an explicit model of the *process* of language comprehension; rather, it is a computational model of pragmatic competence under the cooperative principle (Grice, 1975). In simplest form, a pragmatic listener (L_1) is interested in learning about the world x upon hearing an utterance u.

$$L_1(x \mid u) \propto S_1(u \mid x) \cdot P(x)$$

She does this by integrating her prior beliefs about the world $P(x)$ with the likelihood that a speaker (S_1) would have produced the utterance u given a world x: $S(u \mid x)$. Prior beliefs about the world $P(x)$ capture relevant world-knowledge, which is central to deriving context-sensitive interpretations.

The utterance production model S_1 is a rational (speech-)actor model. According to Bayesian decision theory, the speaker takes actions (i.e., produces utterances) soft-max optimally in accord with his utility function (Baker, Saxe, & Tenenbaum, 2009)[10]. The utility function of a speaker is based on the informativeness of the utterance. Informativeness is computed with respect to what a naive listener L_0 would believe upon hearing the utterance. S_1 thus seeks to minimize the surprisal of s given u for the naive listener, while bearing in the mind the utterance cost $C(u)$.

$$U_{S_1}(u;s) = \ln(L_0(s \mid u)) - C(u)$$

[10]See `agentmodels.org` for more on this form of model specification.

The speaker is then a soft-max optimal agent, with degree of optimality governed by λ.

$$S_1(u \mid x) \propto \exp(\lambda \cdot U_{S_1}(u; s))$$

The speaker's decision is made with respect to a hypothetical literal listener L_0. The literal listener evaluates an utterance according to what the words mean: their truth-functional denotation, or semantics. The semantics is a context-invariant mapping, from utterances and worlds to truth values: $[\![u]\!] : X \rightarrow$ Boolean. From this, we construct a base case of a *literal listener*, who interprets utterances with respect to their literal meaning.

$$L_0(x \mid u) \propto \delta_{[\![u]\!](x)} \cdot P(x)$$

Note that this model is the same as L_1 except for the likelihood function. Here, the likelihood $\delta_{[\![u]\!](x)}$ is the Kronecker delta function returning 1 for states x compatible with utterance u, and 0 otherwise.

The RSA framework formalizes the Gricean maxim to be informative: The speaker produces utterances not with respect to their literal meaning, but with respect to how a naive listener (the literal listener) would interpret them, given the literal meaning and the listener's prior beliefs. The fact that the literal listener updates her beliefs based on the truth-functional meaning of the utterance results in the speaker behaving in accord with the Gricean maxim to be truthful: the speaker who reasons about the literal listener has the goal of communicating the state x given an utterance u and given its literal meaning $[\![u]\!]$. In principle, the speaker–listener recursion could grow indefinitely. In practice, stable pragmatic phenomena can be observed with just one level of recursion.

Table 2: Stimuli used in Experiment 1. Estimates are proportion agreement for truth judgments and Maximum A-Posteriori (MAP) estimates for prevalence. Brackets denote 95% confidence intervals for truth judgments and 95% credible intervals for prevalences.

Conceptual type	Item	Truth judgment	Prevalence
Majority characteristic	1. Leopards have spots.	0.956 [0.912, 0.989]	92.7 [85.9, 99.0]
	2. Ducks have wings.	0.945 [0.89, 0.989]	98.5 [95.2, 99.9]
	3. Cardinals are red.	0.989 [0.967, 1]	75.5 [62.1, 86.9]
	4. Swans are white.	0.901 [0.835, 0.967]	67.3 [57.1, 73.8]
	5. Peacocks have beautiful feathers.	0.989 [0.967, 1]	92.0 [77.7, 100]
Minority characteristic	6. Lions have manes.	0.945 [0.89, 0.989]	54.3 [48.1, 63.0]
	7. Kangaroos have pouches.	0.967 [0.923, 1]	69.9 [65.8, 79.6]
	8. Robins lay eggs.	0.934 [0.879, 0.978]	68.5 [61.1, 74.9]
Striking	9. Sharks attack swimmers.	0.879 [0.813. 0.945]	41.8 [32.0, 54.7]
	10. Mosquitos carry malaria.	0.989 [0.967, 1]	47.2 [38.1, 52.9]
	11. Ticks carry Lyme disease.	0.967 [0.923, 1]	42.6 [40.0, 54.2]
	12. Tigers eat people.	0.692 [0.593, 0.78]	37.3 [23.3, 49.9]
False generalization	13. Robins are female.	0.429 [0.33, 0.527]	51.9 [47.8. 55.4]
	14. Lions are male.	0.571 [0.473, 0.67]	49.9 [46.6, 53.9]
	15. Swans are full-grown.	0.725 [0.637, 0.802]	59.7[51.0, 66.0]
	16. Leopards are juvenile.	0.143 [0.077,0.22]	28.3 [22.9, 38.2]
	17. Sharks are white.	0.341 [0.253, 0.44]	32.2 [15.8, 44.6]
False or Uncertain	18. Leopards have wings.	0.022 [0, 0.055]	1.0 [0.0, 4.1]
	19. Kangaroos have spots.	0.033 [0, 0.077]	5.0 [0.1, 13.5]
	20. Tigers have pouches.	0.011 [0, 0.033]	2.0 [0.0, 12.3]
	21. Robins carry malaria.	0.055 [0.011, 0.099]	5.4 [1.6, 9.1]
	22. Sharks have manes.	0.044 [0.011, 0.099]	0.3 [0.0, 5.7]
	23. Lions lay eggs.	0 [0, 0]	0.1 [0.0, 4.3]
	24. Sharks don't attack swimmers.	0.231 [0.154, 0.319]	59.3 [49.8, 75.1]
	25. Ticks don't carry Lyme disease.	0.044 [0.011, 0.088]	55.1 [46.6, 60.8]]
	26. Mosquitos don't carry malaria.	0.077 [0.033, 0.132]	58.7 [50.5, 65.3]]
	27. Tigers don't eat people.	0.297 [0.198, 0.385]	65.3 [56.9, 97.7]
	28. Peacocks don't have beautiful feathers.	0.022 [0, 0.055]	17.4 [0.5, 27.5]
	29. Mosquitos attack swimmers.	0.385 [0.286, 0.495]	26.5 [8.4, 39.2]
	30. Sharks lay eggs.	0.088 [0.033, 0.143]	18.5 [0.6, 41.3]

Table 3: Stimuli used in Experiment 2 and statistics of the priors measure in the prior elicitation task. *Potential to be present* is a measure of how many different kinds are expected to have the property. Mean prevalence when present is a measure of how widespread the property is expected to be, assuming that it is present within a kind. Maximum A-Posteriori (MAP) estimates and 95% Highest Probability Density Intervals are shown for each measure. The majority of these items are taken from Cimpian et al. (2010)

Property type	Item	Potential to be present θ	Mean prevalence when present γ
Body part	fur	0.78 [0.81, 0.69]	0.90 [0.828, 0.93]
	skin	0.89 [0.93, 0.85]	0.89 [0.853, 0.93]
	feathers	0.68 [0.74, 0.61]	0.90 [0.827, 0.91]
	legs	0.88 [0.93, 0.81]	0.87 [0.819, 0.94]
	tails	0.75 [0.82, 0.65]	0.91 [0.881, 0.93]
	ears	0.85 [0.89, 0.80]	0.88 [0.839, 0.91]
	claws	0.73 [0.74, 0.61]	0.87 [0.749, 0.89]
	teeth	0.83 [0.88, 0.77]	0.88 [0.810, 0.91]
Color	silver legs	0.37 [0.46, 0.30]	0.58 [0.472, 0.64]
	yellow fur	0.53 [0.65, 0.48]	0.69 [0.594, 0.76]
	violet skin	0.39 [0.51, 0.33]	0.63 [0.524, 0.74]
	orange ears	0.46 [0.54, 0.35]	0.60 [0.488, 0.69]
	blue claws	0.38 [0.44, 0.26]	0.66 [0.481, 0.69]
	pink teeth	0.22 [0.37, 0.21]	0.54 [0.421, 0.66]
	orange tails	0.51 [0.57, 0.38]	0.65 [0.581, 0.75]
	purple feathers	0.48 [0.54, 0.36]	0.64 [0.509, 0.72]
Vague	big claws	0.63 [0.70, 0.55]	0.78 [0.682, 0.84]
	long teeth	0.60 [0.66, 0.50]	0.73 [0.694, 0.84]
	rough skin	0.62 [0.72, 0.55]	0.74 [0.672, 0.80]
	curly fur	0.55 [0.64, 0.48]	0.76 [0.669, 0.82]
	long legs	0.63 [0.68, 0.55]	0.76 [0.680, 0.83]
	smooth feathers	0.62 [0.68, 0.50]	0.80 [0.692, 0.83]
	long tails	0.62 [0.69, 0.53]	0.82 [0.738, 0.85]
	small ears	0.65 [0.70, 0.55]	0.81 [0.751, 0.85]
Common accidental	torn tails	0.47 [0.54, 0.32]	0.23 [0.202, 0.35]
	wet fur	0.57 [0.66, 0.47]	0.45 [0.376, 0.56]
	dusty skin	0.45 [0.56, 0.39]	0.44 [0.393, 0.58]
	torn feathers	0.46 [0.58, 0.40]	0.29 [0.218, 0.35]
	fungus-covered fur	0.37 [0.51, 0.29]	0.37 [0.294, 0.49]
	worn-out claws	0.54 [0.62, 0.46]	0.41 [0.336, 0.54]
	muddy feathers	0.41 [0.58, 0.36]	0.39 [0.291, 0.45]
	sore teeth	0.39 [0.51, 0.33]	0.25 [0.188, 0.33]
Rare accidental	broken legs	0.30 [0.44, 0.23]	0.13 [0.098, 0.17]
	swollen ears	0.42 [0.52, 0.32]	0.24 [0.196, 0.34]
	itchy tails	0.37 [0.48, 0.28]	0.33 [0.231, 0.38]
	rotten teeth	0.56 [0.62, 0.42]	0.31 [0.225, 0.39]
	sore legs	0.43 [0.53, 0.34]	0.24 [0.209, 0.34]
	cracked claws	0.50 [0.60, 0.36]	0.24 [0.194, 0.36]
	infected ears	0.42 [0.54, 0.34]	0.18 [0.131, 0.23]
	burned skin	0.32 [0.40, 0.23]	0.19 [0.133, 0.25]